



# Fast Data Loading into the Vertica<sup>®</sup> Analytic Database

*Vertica data loading architecture, benchmarks and best practices*

## Fast Data Loading into Vertica® Analytic Database

The grid-based, column-oriented Vertica® Analytic Database is well known for its blinding query performance against large volumes of data, but unlike other column databases, Vertica has the unique ability to insert data at very high rates of speed. It can even load on a non-stop basis while the database is being queried, thus enabling real-time analysis. This paper describes how Vertica loads data and also highlights customer cases to illustrate loading speeds achieved in production.

### Loading Methods




There are three methods used to load data into a Vertica database, all of which can run while the database is concurrently queried:

1. **Trickle Load** – insert data into Vertica as it arrives
2. **Bulk Load to Memory** – ideal for fast loading of smaller data volumes (a few gigabytes)
3. **Bulk Load to Disk** – ideal for large data volumes (many giga- or terabytes)

Vertica is capable of loading many megabytes/second per machine per load stream. With Vertica’s shared-nothing architecture, loading can be performed in parallel, scaling linearly as servers are added to the Vertica cluster.

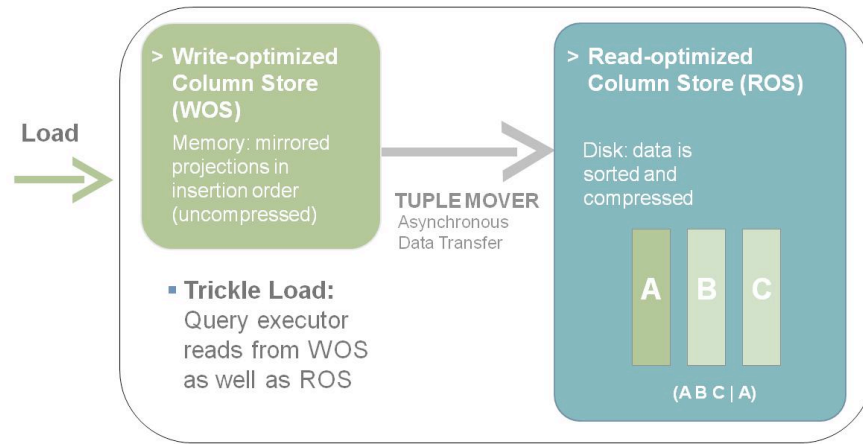
### Vertica Customer Loading Experiences

Before examining how Vertica loads data this rapidly, here is a look at some load speeds reported by Vertica customers:

|   |  |
|---|--|
|  | <p><b>Trickle loads 48MB/minute</b> – SNMP data generated by devices in the Comcast cable network is trickle loaded on a 24x7 basis at rates as high as 135,000 rows/second. The system runs on 5 HP ProLiant DL 380 servers.</p>                          |
|  | <p><b>Bulk loads to memory 300MB/minute</b> – 50MB to 300MB of call detail records (1K record size—150 columns per row) are loaded every 10 minutes. Vertica runs on 6 HP ProLiant DL380 servers.</p>  |
|  | <p><b>Bulk loads to disk 5GB/minute</b> - The loading and enrichment (i.e., summary table creation) of 1.5TB of call detail records formerly took 5 days in a row-oriented data warehouse database. Vertica required 5 hours to load the same data.</p>    |
| <p><b>Global investment firm</b></p>  | <p><b>Trickle loads 2.6GB/minute</b> - Historic financial trade and quote (TaQ) data was bulk loaded into the database at a rate of 125GB/hour. New TaQ data was trickled into the database at rates as high as 90,000 rows per second (480b per row).</p> |

## Why Vertica Loads Data Quickly

Vertica achieves fast loading by cleverly utilizing main-memory caching, batch processing and parallelization. Firstly, Vertica caches new, incoming data into a main-memory-based Write Optimized Store (WOS). In the WOS, the data is stored in columns and can be queried, but the data is neither compressed nor sorted yet. Then, an asynchronous process called the Tuple Mover, moves batches of data out of the WOS and into



disk-based Read Optimized Store (ROS) containers, compressing and sorting the columns in the process. In this way, the performance cost of sorting and compression is shared and amortized across many tuples. The expected number of disk writes per INSERT will also decrease as the batch size grows. Thirdly, since Vertica database tables are partitioned across a collection of shared-nothing nodes in a grid, it distributes the INSERTs evenly across the cluster based on a partitioning key. If the INSERT rate grows, more nodes can be added to cope with the increase.

A fair question to ask is whether this approach would mean that the answers to queries would be stale. The answer is no ... the Vertica query evaluator searches both places (the WOS and the ROS) for data to answer queries.

## A Note about ACID

Of course, in order to support ACID transactions in this setting, there must be a safe way to allow committed data to reside in main memory. Vertica accomplishes this by keeping redundant copies of the data in multiple WOS/ROSeS on different nodes. Vertica achieves k-safety, where k is the number of nodes that can fail without losing any work, by keeping data copies on k+1 different machines. All INSERTs will be sent to all k+1 relevant sites and stored in their main memory caches. Once all these copies are installed, the tuple is stable (subject to the k-safety constraints).

## Loading Techniques

Data can be loaded into Vertica using either the SQL INSERT or the COPY command. SQL INSERT “trickle” loads data needed for real-time analysis into the WOS.

The SQL COPY command loads data from a file or input stream (generated by ETL software or another data feed), where the data can be located local or remotely and may also be stored in a compressed format like GZIP. For small batches of data (a few GB) which must be loaded very quickly, data is loaded into the WOS. When loading large quantities of data, the recommended approach is to use the COPY command with the DIRECT option. Data is then loaded directly into the ROS on disk which is the fastest way to load large batches of data. Since COPY accepts a delimited file, it is very easy to output your data from the source system, into a format which Vertica understands and can quickly load. A typical COPY command is shown below which loads data into the fact table from data file fact.txt, which is delimited with a “|,” and all exceptions and rejected records are written into the specified log, and the DIRECT option loads data directly into the ROS.

```
COPY fact FROM 'fact.txt' DELIMITER '|'
EXCEPTIONS 'except.log' REJECTED DATA 'rejects.log'
DIRECT NULL '\\';
```

## Tips for Faster Loading

Loading data into the Vertica Analytic Database is very straightforward; here are some additional tips to ensure that loads run as fast as possible:

### Bulk Loading

1. Use COPY with DIRECT so that the data is written directly to the disk.
2. Load data in large chunks and if necessary merge files together as one. Monitor this by checking how many ROS storage containers exist.
3. Via an application, use two buffers, where one is buffering and the other is COPYing, or...stream data into a COPY statement using multiple concurrent streams.
4. Ensure maximum memory available.
5. Load multiple streams on different nodes.

### Trickle Loading

1. Use COPY whenever possible.
2. Use INSERT when loading < 1000 rows, otherwise use COPY.
3. Load multiple streams on different nodes.
4. Ensure maximum memory available.

## See for yourself...

The Vertica Analytic Database supports SQL and integrates with 3<sup>rd</sup>-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 Database or if you would like to evaluate it yourself, then visit the following links:

|                                     |  |  |
|-------------------------------------|--|--|
| <b>Vertica Resource Library</b>     | <a href="http://www.vertica.com/resourcelibrary">www.vertica.com/resourcelibrary</a> | White papers, demos, webcasts, system requirements   |
| <b>Vertica Benchmarks</b>           | <a href="http://www.vertica.com/benchmarks">www.vertica.com/benchmarks</a>           | See customer-submitted cost and performance comparisons between Vertica and other databases                    |
| <b>Vertica for the Cloud</b>        | <a href="http://www.vertica.com/cloud">www.vertica.com/cloud</a>                     | Get a Vertica database instance provisioned instantly on the Amazon Cloud and use it on a month-to-month basis |
| <b>Vertica Customers</b>            | <a href="http://www.vertica.com/customers">www.vertica.com/customers</a>             | See who's using Vertica  |
| <b>Request a Vertica Evaluation</b> | <a href="http://www.vertica.com/download">www.vertica.com/download</a>               | 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>.