



# Vertica

# Review From A Customer

From IT Central Station, the leading review site for enterprise technology solutions.

---

## Review by a Real User

Verified by IT Central Station



Staff Dev Lead - Analytics Data Storage at a tech services company with 1,001-5,000 employees

it\_user692295

### WHAT IS MOST VALUABLE?

Two of them: The core feature, meaning their highly efficient columnar file format and execution engine along with a great coverage of ANSI SQL, provides our analysts with a highly expressive and performing platform. The extensibility and efficiency provided by their C++ SDK.

### HOW HAS IT HELPED MY ORGANIZATION?

Before Vertica, we used a combination of sharded RDBMSs and Hive: the typical runtime for a query was in the hours. It's now in the seconds, with waymore data than then (we're talking hundreds of terabytes).

### WHAT NEEDS IMPROVEMENT?

Whatever's out, the core is not always as great as the engine, especially their first version. That's true, for example, for the Kafka or Hadoop integration. But they're getting better release after release.

### FOR HOW LONG HAVE I USED THE SOLUTION?

Four years.

### WHAT DO I THINK ABOUT THE STABILITY OF THE SOLUTION?

Vertica's code, being designed to use the hardware at its maximum, is very sensitive to low level changes such as kernel bumps or GLibC upgrades. It's also important to do tests on the storage layer (RAID controller + disks).

### WHAT DO I THINK ABOUT THE SCALABILITY OF THE SOLUTION?

The default layout (all nodes running spread) introduces latencies in query planning when you reach about 60 nodes, in our experience. Switching to a large cluster (one control node per rack) would be advised, way before reaching the 128 nodes hard limit.

#### **HOW ARE CUSTOMER SERVICE AND TECHNICAL SUPPORT?**

It's really great. One of the best I had to deal with. They also assisted us during the development phase of the custom components we've designed.

#### **WHICH SOLUTION DID I USE PREVIOUSLY AND WHY DID I SWITCH?**

Not really in the same area (MPP databases). However, we ran benchmarks back then against a bunch of competitors and Vertica was one of the fastest, while being relatively cheap and able to accommodate our hardware.

#### **HOW WAS THE INITIAL SETUP?**

The setup per se was pretty straightforward. However, it took us some time to design the most efficient loading pattern from Hadoop.

#### **WHAT'S MY EXPERIENCE WITH PRICING, SETUP COST, AND LICENSING?**

Nothing to advise really; try it out first, it's free up to three nodes and 1TB, and then get in contact with their sales guys.

#### **WHICH OTHER SOLUTIONS DID I EVALUATE?**

We did evaluate mostly SAP HANA and SQL Server PDW back in 2013, along with a bunch of OSS solutions.

#### **WHAT OTHER ADVICE DO I HAVE?**

If you plan to use Vertica for different workloads (in term of IO patterns, query frequency, dataset structure) plan to split your clusters: the mother of all cluster patterns becomes quite difficult to manage at some point. We today have around 20 clusters for different usages.

Learn more: [Read 10 reviews of Vertica](#)