

Vertica Accelerator

Vertica's Database-as-a-Service (DBaaS) offering delivers a unified, high-performance, advanced analytics and machine learning platform with automated cloud setup, administration, and management on Amazon Web Services.

Product Overview

Available as software as a service (SaaS) on AWS, Vertica Accelerator provides all the functionality of the Vertica analytical database. It runs in your own AWS cloud account, and it is managed by Vertica. All of your data analysts and data scientists can tap the power of Vertica advanced analytics and machine learning without the logistical issues associated with scaling, upgrades, monitoring, or support.

Vertica Unified Analytics Platform

The Vertica Unified Analytics Platform provides blazingly fast speed (queries run 10–50x faster), exabyte scale with aggressive compression (store 10–30x more data), openness (use any BI/ETL tools, storage formats, etc.) all at a much lower cost than traditional data warehouse solutions. Proven over years of deployment with hundreds of customers in a wide variety of industries, cloud-native Vertica uses cloud object storage and separately scaled compute to deliver on the promise of self-service analytics for the many types of analytical users in your business. Our development strategy has always been about deploying anywhere, from bare metal servers to data warehouse as a service, or from internally managed servers to self-service cloud.

Key Benefits

What Are the Advantages?

If you are considering moving analytic workloads to the cloud, every company needs the following:

- Analytics that scales: As your data and

analysis needs grow – more users, more use cases, more data – Vertica Accelerator grows with you.

- **Faster analysis:** Finish queries in minutes, not hours; seconds, not minutes. Accomplish analytical use cases you once thought impossible.
- **Easy administration:** Vertica handles upgrades, backups, and installation.
- **Data democratization:** Provides database access to all analysts and data scientists.
- **Workload isolation** ensures separate teams can work without conflicts or slowdowns.
- **No extra charges:** Move from traditional BI to advanced analytics like geospatial, time series, or machine learning without paying extra.
- **Advanced analytics:** Data preparation for data science, feature correlation and selection, time series analysis, event pattern matching, geospatial analysis, machine learning algorithms, model evaluation and management – all built-in. More than 600 analytics functions.
- **Freedom to negotiate:** Use your negotiated discounts, savings plans, and reserved instances set up on your secure AWS account.

Vertica Accelerator offers these features and more.

Key Features

At the core of Vertica Accelerator is the same column-oriented, relational database that Vertica customers trust, built to handle today's analytic workloads. Vertica Accelerator provides:

- **Complete and advanced SQL-based analytical functions** to provide powerful SQL analytics.
- **Extreme compression on large data sets**, offering superior query and analytic performance.
- **Full use of data in your own AWS account** –



you can reduce risk and comply with regulations, without metadata sharing.

- **Autoscaling** to ramp up compute automatically, but with user-defined guardrails to avoid the surprise of high end-of-month bills.
- **High concurrency** with isolated workloads so reporting, BI, ML, ETL, and dozens of users can work simultaneously.
- **In-database machine learning algorithms** and R, Python, PMML, TensorFlow extensibility.
- **Easy administration** with upgrades, backups, and installation – all handled for you.
- **Availability Zone (AZ) failover and node recovery** – Spend more time getting value from your data and less time with administration as you scale to more use cases and higher data volumes with fewer resources.

Contact us at
www.vertica.com

Like what you read? Share it.



How Does It Work?

To get started, visit vertica.com/accelerator. Sign up for a free trial of the service. Vertica experts ensure you configure the most efficient and cost-optimized database. The migration tooling automates data loading and setup. Vertica implements cross-account access and proper security protocols for you to then start analyzing.

Analyze More for Less

Spend Less on Compute

Vertica has a long history of optimizing hardware utilization – we use fewer nodes to accomplish the same thing faster. Many cloud-based databases re-sell hardware at a markup. The more you use, the more you pay them, so they have no financial incentive to reduce either compute or storage.

Spend Less on Storage

Aggressive data compression (3-30x more than competitors) means you pay less to store all of your data.

Autoscaling

With No Bill Shock

Like most data warehouse as a service options, Vertica Accelerator autoscales compute to handle varying demands, but it also lets you define

guardrails to avoid the surprise of enormous cloud bills. In addition, auto scheduling allows you to schedule larger compute workloads for times like end of quarter when you know your work demand is higher.

Keep Your Data in Your Own Cloud

Keep Data Secure

Multi-tenant cloud implementations often store metadata for multiple customers in a single database, creating a security issue. There is also the performance challenge caused by shared infrastructure. One customer may use a great deal of resources, reducing the available resources for other clients. This “noisy neighbor effect” means your performance may suffer even when you are not overusing the database.

Vertica Accelerator is single-tenant. You own all your own bandwidth, so there is no noisy neighbor effect, and all metadata is stored in your own instance—not in a shared system. Keep your data secure in your own AWS account, with your own organization’s security rules in place.

concurrent users. Plus, streaming data demands that ingestion, transformation, and data preparation be fast and constant. Reporting, machine learning, geospatial, ad-hoc queries—all tasks need to run simultaneously, without bogging down performance. That is why Vertica’s sub-clusters isolate workloads so you can have as many simultaneous users doing as many analytics tasks as you wish without slowing each other.

SQL, Python, R: It’s Up to You

Use Most Productive Interface

Every function is available in SQL, making integration with visualization tools a snap. But if Python is more comfortable for you, you can use the open source library VerticaPy and put Vertica’s powerful database engine to work analyzing whole data sets from a Jupyter notebook. Import/export PMML or train TensorFlow graphs, import them, and deploy models in minutes, to put advanced analytics to work.

Concurrency

How do I do ALL the things I need?

The business imperative to become more data-driven builds the need for more and more

