



Rocket® Vertica®

A deploy-anywhere SQL database engineered for massive scale, elasticity, and advanced analytics.

At a glance

- Native support for SQL, Python, and R to democratize analytics across your team.
- Flexible deployment across SaaS, public cloud, & on-premises environments.
- Direct queries on your existing data lakes without costly or risky data relocation.
- Robust, format-preserving encryption for maximum data security and compliance.

Modernizing your data warehouse shouldn't mean compromising on speed or scale. Rocket® Vertica® provides an open, elastic Vertica database built to handle your largest workloads. By bringing advanced analytics and machine learning directly to your data, we help you uncover critical trends, reduce infrastructure overhead, and confidently build the foundation for your data-driven future.

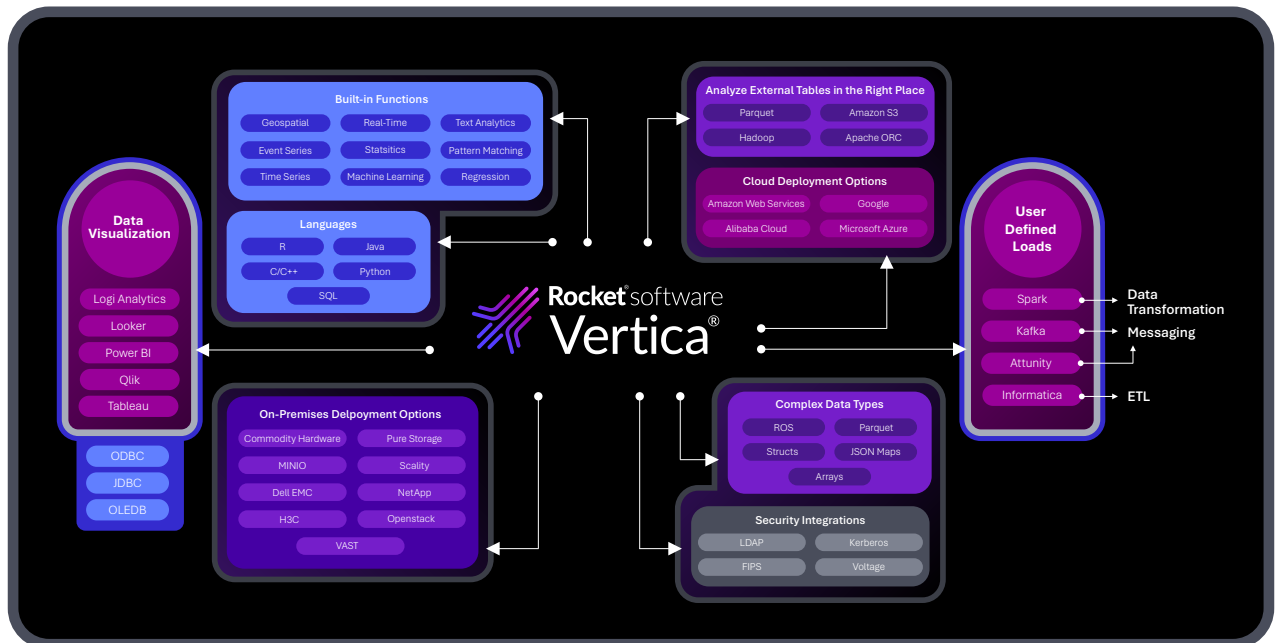


Fig 1: Rocket Vertica's open architecture and rich ecosystem

Key benefits

- 01 Cloud-native elasticity**
Eon Mode separates compute and storage, allowing you to scale Vertica resources independently and optimize your operational costs.
- 02 Lightning-fast performance**
Massively parallel processing powers through complex queries in seconds, eliminating the frustrating bottlenecks of traditional row-store databases.
- 03 In-database machine learning**
Train, build, and deploy predictive AI models directly inside your secure Vertica data cluster.
- 04 Maximized storage efficiency**
Advanced compression capabilities let you store up to 30x more data using significantly fewer servers.
- 05 Complete ecosystem integration**
Preserve your investments by seamlessly connecting to the BI, ETL, and visualization tools you already use.
- 06 True hybrid flexibility**
Move workloads freely between multi-cloud and on-premises setups to avoid vendor lock-in and maximize your existing infrastructure.
- 07 Enterprise-grade security**
Protect data at rest and in motion with fine-grained access controls, TLS, and LDAP integration.

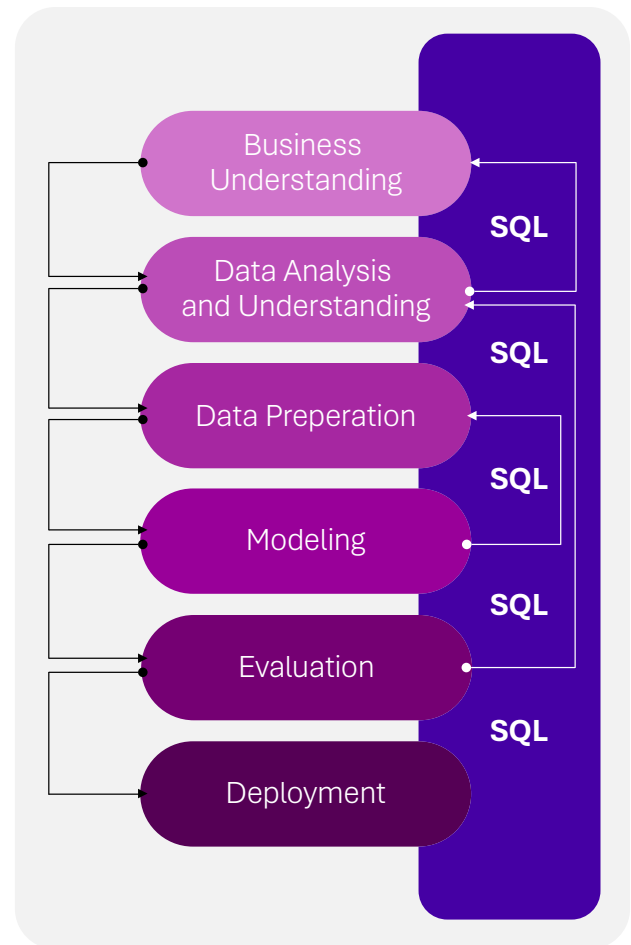


Fig 2: Rocket Vertica built-in machine learning process flow



With Rocket Vertica, you can:

01

Analyze massive datasets directly in your data lake without moving them, reducing time to insight and system overhead.

02

Scale compute power dynamically to support peak workload demands without overprovisioning your resources.

03

Run complex predictive analytics and machine learning algorithms using the native programming tools your data scientists prefer.

04

Prevent resource competition between departments by securely isolating specific BI and data science workloads with sub-clustering.

05

Avoid cloud vendor lock-in by maintaining complete workload portability across public clouds and private data center environments.

06

Maintain strict regulatory compliance with granular masking and data-centric protection for all your sensitive customer information

Insights at enterprise scale

- Over 700 built-in analytical functions
- 10-50X faster query performance
- Store 10-30X more data per server
- 100% separation of compute and storage

Accelerate your journey to data-driven success

[Learn more](#)

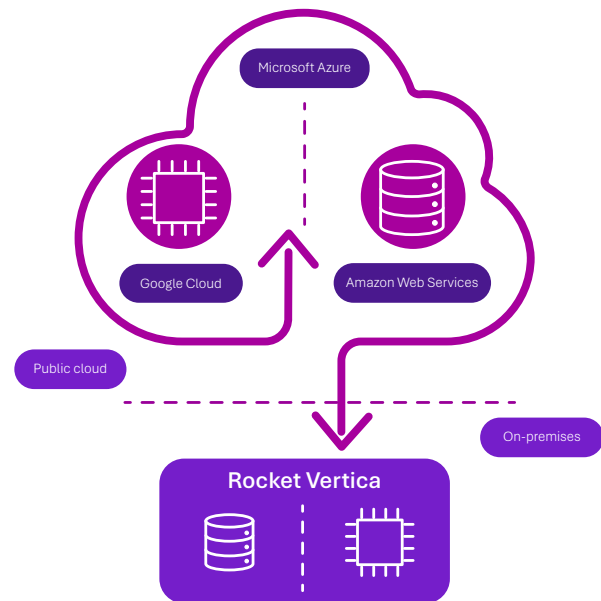


Fig 3: Analytics powered by Rocket Vertica across on-premises and public cloud