



Get the value you were promised from your Hadoop investment

Analyzing data from across the business makes sense both for business intelligence and data science, but Hadoop hasn't lived up to the hype. Get concurrent workloads, fast query speed, and big ROI by adding Vertica in Eon Mode for HDFS.

Vertica in Eon Mode for HDFS:

Operational Flexibility:

- Leverage the separation of compute and storage architecture to scale compute resources up or down easily based on demand.

Workload Isolation:

- Protect workloads of departments, teams, or projects from cross-interference without replicating data.

Benefits of Cloud Innovation:

- Vertica in Eon Mode for HDFS delivers the benefits of cloud data centers, without the hidden risks of a massive migration project to get your data and analytics to the cloud.

Future-Proof Analytics:

- As business needs change, move to any major cloud, to a hybrid model, or simply expand on-premises with the same blazing fast analytics.

Hadoop Hope and Hype

A huge number of businesses have invested in Hadoop clusters, data lakes that were supposed to unite their data into one place, and provide tremendous returns with unlimited business intelligence, and new, even better forms of analyses like machine learning.

As the hype around Hadoop has fallen away, many companies wonder if the money they invested was wasted, or if there is a way to get the unified analytics they hoped for.

The answer is yes. There is a way, but you're not going to find it in a Hadoop distribution. Vast quantities of data from useful sources can reveal lucrative patterns that make that data collection worthwhile, and the Hadoop Distributed File System (HDFS) is an excellent, scalable, affordable, and resilient way to store all that data. Keep the data and the storage method, just change the way you analyze that data.

Query Engines Don't Cut It

Hadoop query engines such as Impala, Hive, and Presto have been the biggest disappointments of the Hadoop ecosystem. Companies keep seeing their analytics work on a small scale, then fail when they try to do the same thing in production with a hundred users, or a half dozen different teams. Organizations see as many as 25% of their daily queries fail.

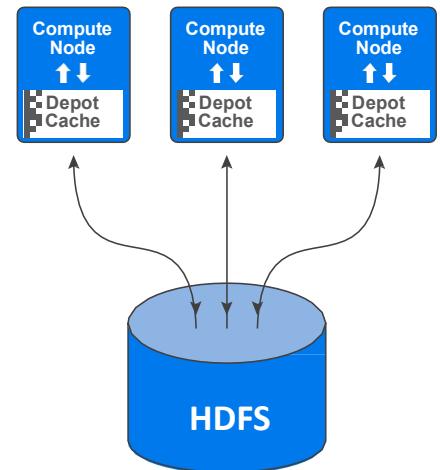


Figure 1. Vertica in Eon Mode for HDFS stores persistent data in HDFS. Compute nodes fill local cache from HDFS for fast queries. This allows organizations to isolate workloads, speed queries, and provide analytic access to many different teams.

Vertica + Hadoop = Analytics ROI

Vertica offers a wide variety of options to get analytic value out of Hadoop. Open API's and a powerful ability to integrate with Hadoop file formats and ecosystem products provide that flexibility. While many vendors pursue the traditional concept that you must put all your data in one place, Vertica understands that it is the analytics you need to unify, not the data.

Contact us at:
www.vertica.com

Like what you read? Share it.



Unified Analytics Warehouse

The data warehouse has been an analytics workhorse for decades. Unprecedented volumes of data, new types of data, and the need for advanced analyses like machine learning brought on the age of the data lake. But Hadoop by itself can't deliver the value businesses need. Now, many companies have a data lake, a data warehouse, or a mishmash of both, possibly combined with a mandate to go to the cloud. The end result can be a sprawling mess, a lot of duplicated effort, a lot of missed opportunities, a lot of projects that never made it into production, and a lot of financial investment without return.

The Unified Analytics Warehouse combines the best of both worlds, and throws in some of the flexible scaling of the cloud. Technical and spiritual unification of the two opposed camps makes a powerful impact on the effectiveness of analytics for the business over all.

Learn more at
<https://www.vertica.com/vertica-for-hadoop/>

Vertica in Eon Mode for HDFS

Vertica in Eon Mode for HDFS delivers analytical database power with a separation of compute and storage architecture similar to most cloud architectures, delivered on-premises.

- Isolate workloads by business unit or by type of workload using subclusters.
- Scale compute to the active data set ("hot data") without archiving inactive data.
- Bring varying levels of compute to the database by scaling the compute nodes as needed for seasonal workload peaks.
- Hibernate compute completely when not needed, freeing up compute for other purposes. Then revive back into nodes when needed again.
- Rebalance compute nodes on the fly. Get high availability with simplified database management.

Analyze Data Wherever It Lives

While the public cloud is a necessary consideration for most organizations, the future of infrastructure is changeable. Keep your options open for multi-cloud and hybrid — a mixture of on-premises and cloud environments. Not every database or workload belongs in the cloud, but that doesn't mean they can't benefit from cloud innovations. Vertica in Eon Mode for HDFS delivers a scalable, performant analytical data warehouse solution with many of the benefits of the cloud, right on your Hadoop cluster.

Isolate Workloads

The Unified Analytics Warehouse allows you to do multiple analytic workloads without impacting each other from a performance or availability perspective. Run ad hoc business intelligence queries, dashboard reports, and machine learning simultaneously, all in the same database cluster, without resource contention.

Unified Analytics Warehouse

