



**INDUSTRY**  
Data virtualization

**WEBSITE**  
www.denodo.com

**COMPANY OVERVIEW**  
Denodo is the leader in data virtualization software that enables businesses with timely and comprehensive strategic business information.

Headquartered in Palo Alto, CA and with offices and a partner network worldwide, Denodo provides a data integration and data virtualization software called the Denodo Platform for data virtualization, as well as support, training, and consulting services.

**PRODUCT OVERVIEW**  
The Denodo Platform offers the broadest access to structured and unstructured data residing in enterprise, big data, and cloud sources, in both batch and real-time, exceeding the performance needs of data-intensive organizations for both analytical and operational use cases, delivered in a much shorter timeframe than traditional data integration tools.

# HPE and Denodo

## Denodo Expands the Logical Data Warehouse with Support for HPE Vertica

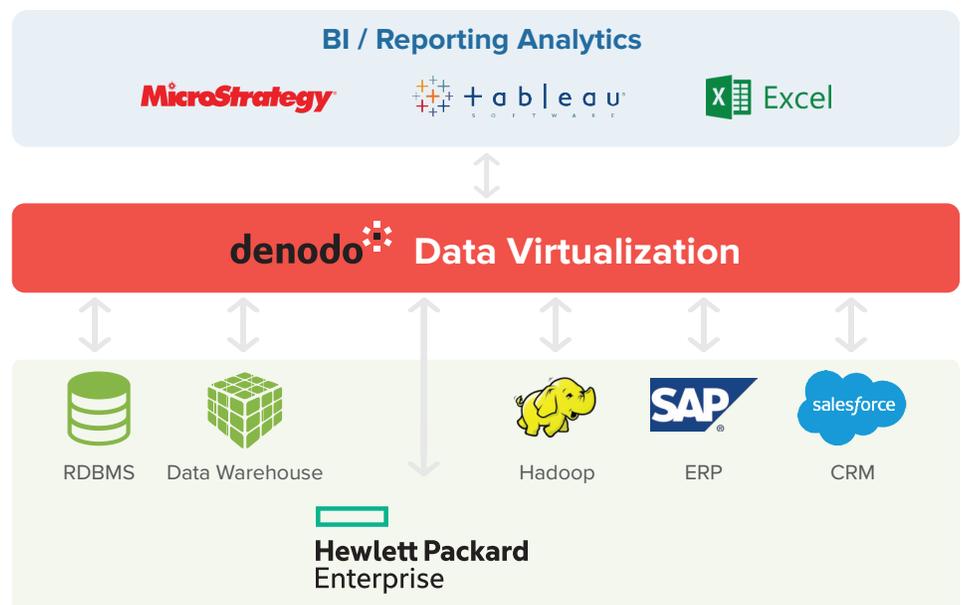
The **Denodo Platform**, with its data abstraction layer, enables the logical data warehouse, which integrates traditional data warehouses with other analytical sources, extending their capabilities. Such sources include analytical appliances, Hadoop clusters, NoSQL databases, and SaaS and other enterprise applications, decoupling reporting tools from data stores, for greater business agility. Denodo has achieved certification in Hewlett Packard Enterprise (HPE) Vertica, enabling Denodo users to integrate HPE Vertica within their logical data warehouse.

## The Denodo Platform Enables the Logical Data Warehouse

Mark Beyer of Gartner introduced the term “logical data warehouse” back in 2012, defining it as “a new data management architecture for analytics, which combines the strengths of traditional repository warehouses with alternative data management and access strategies.” Gartner predicted that the logical data warehouse would form a new best practice by the end of 2015.

Data virtualization is the foundation for the logical data warehouse, as it adds a data abstraction layer that decouples the BI tools from the different data stores (enterprise data warehouses, databases including NoSQL databases, big data stores, etc.).

Data virtualization enables queries to be federated across multiple data sources – both traditional structured data sources, such as databases and data warehouses, and less traditional ones such as Hadoop clusters, NoSQL databases, web services, SaaS applications – while still appearing as a single “logical” data source to the user.



# The Benefits of Using Denodo with HPE Vertica



**Information agility:** The Denodo Platform, when used with HPE Vertica, enables users to generate reports much more quickly, without their having to maintain complex scripts for carrying out data replication across different data stores.



**Advanced performance:** Working together, the Denodo Platform and HPE Vertica optimize performance levels. Denodo optimizes query processing in the logical data warehouse, exploiting HPE Vertica's powerful execution capabilities.



**Scalability:** The Denodo Platform makes it easy to add new data stores to the logical data warehouse whenever they are needed, and it also distributes the processing among HPE Vertica and other stores.



**Hybrid data:** Hybrid data patterns such as data warehouse offloading (DWOL) can be built easily by combining the Denodo Platform, HPE Vertica, and Hadoop. Hot data that needs blazing-fast processing can reside in HPE Vertica, while cold data can be stored in Hadoop. The Denodo Platform would prudently delegate queries to both repositories, achieving maximum performance.

## Benefits of the Logical Data Warehouse

- **Abstraction/decoupling:** As the data model is defined in the data virtualization layer and is exposed to BI tools for consumption, any change in the underlying stores is buffered in this layer without impacting the reporting layer.

In a traditional architecture, BI tools are completely tied to specific data stores (e.g. BI + EDW or BI + Hive) and any changes will impact them. Users have to know the details of how the data is stored in each of the different stores. With the logical data warehouse, the user has more flexibility for accommodating changes in the data stores (e.g. adding new fields to a table in an EDW or in Hive, adding new tables, or integrating a new store or technology, e.g., a new NoSQL store, Apache Spark in Hadoop, etc.).

Users don't need to know the details of where the information is stored; they only need to query existing data views published in the data virtualization layer. Users can carry out the metadata changes in the data virtualization layer quickly, without impacting the reporting layer.

The logical data warehouse is a future-proof architecture that allows the user to adopt new technologies with minimum impact on business applications, which is especially important in the big data space.

- **Single definition of data specifications (logical data model):** With the logical data warehouse, there is no need to build a semantic layer in every BI tool, which usually leads to data inconsistencies. All BI tools can share the same model, for improved data quality and fewer data inconsistencies.
- **Reusability:** As a shared data service layer is exposed in terms of canonical views, any business application/reporting tool/web portal can consume the services, improving reusability and reducing cost. In a traditional architecture, models defined in a BI tool are not reusable.
- **Better data governance and security:** This metadata layer can be managed from a single point, from which the user can enforce data governance and security rules. The logical data warehouse provides comprehensive support for data lineage, change impact analysis, the definition of security policies, and so on.
- **Performance:** Data virtualization offers advanced query optimization techniques that are not usually available in BI tools, offering excellent performance.

---

## About Denodo

Denodo is the leader in data virtualization providing agile, high performance data integration and data abstraction across the broadest range of enterprise, cloud, big data and unstructured sources as real-time data services at half the cost of traditional approaches. Denodo's customers across every major industry have gained significant business agility and ROI by enabling faster and easier access to unified business information needs for agile BI, big data analytics, web and cloud integration, single-view applications, and enterprise data services. Denodo is well-funded, profitable and privately held.

[www.denodo.com](http://www.denodo.com)

---

## About HPE

Hewlett Packard Enterprise is an industry leading technology company that enables customers to go further, faster. With one of the industry's most comprehensive portfolios, spanning the cloud to the data center to workplace applications, HPE technology and services help customers around the world make IT more efficient, more productive, and more secure.

[www.hpe.com](http://www.hpe.com)

# High Performance Analytics in the Logical Data Warehouse with HPE Vertica and Denodo

HPE Vertica is an advanced, standards-based relational database for SQL analytics that offers interactive, blazing-fast analytics. Denodo is certified for HPE Vertica, so Denodo users will be able to integrate HPE Vertica into their logical data warehouse architecture.

Using Denodo in conjunction with HPE Vertica enables users to achieve very high performance even when data can be distributed across HPE Vertica and other data stores. The Denodo Platform's Dynamic Query Optimizer applies a cost model that supports HPE's columnar-based approach and applies the push-down optimization technique, which delegates query processing directly to HPE Vertica.

The Denodo Platform makes use of advanced query optimization techniques, including full and partial aggregation pushdown, which rewrites queries so that aggregation operations (e.g. groupby) are moved down to lower nodes in the query execution tree, enabling users to delegate their executions fully or partially into high-performance stores such as HPE Vertica, which can carry out aggregation operations very efficiently. The Denodo Platform's advanced cost model accommodates execution internals of the different data stores, to select the most efficient strategy for each query execution. In the context of this technical certification, Denodo has extended its cost model to include the specifics of the column-based HPE Vertica.

In addition, HPE Vertica can be used as a high-performance cache in the Denodo Platform, enhancing performance for frequently accessed reports.

With the logical data warehouse architecture, the user has more flexibility over where to keep the data, avoiding massive data replication between the different stores. A typical approach is to move fact tables into HPE Vertica and leave dimension tables in other stores (e.g. a product database for the product dimension). However, using Denodo data virtualization, users can build a virtual data mart on top of the different stores and expose a single access point to reporting tools. This offers greater agility to business users, as it enables them to generate reports more quickly and inexpensively.

Finally, Denodo offers advanced capabilities for accessing data from Hadoop, web, cloud, and SaaS sources, and can be used to expand the reach of analytics in HPE Vertica with data from external sources, enhancing business insights.



[denodo.com](http://denodo.com)  
[twitter.com/denodo](https://twitter.com/denodo)



**Hewlett Packard  
Enterprise**

[hpe.com](http://hpe.com)  
[twitter.com/HPE](https://twitter.com/HPE)