



**denodo**  
DATA VIRTUALIZATION

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

525 University Avenue, Suite 31, Palo Alto, CA 94301, USA  
Tel: (+1) 650 566 8833 Email: [info@denodo.com](mailto:info@denodo.com)

## Denodo

### The Company

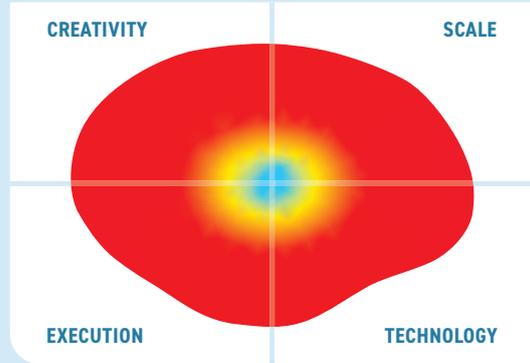
Denodo was founded in 1999 and from its inception has been a pioneer in data virtualisation, which is a key underpinning of the newer architectural concept of a data fabric. It now has over 1,100 customers around the world, over 800 employees and 250 partner companies. Customers include Sanofi, Toyota, ABN-AMRO and BNP Paribas.

### What is it?

In a data fabric, a data catalogue maintains an inventory of enterprise data sources and a semantic layer is constructed that sits above the various source systems. Authorised business users use the semantic layer to identify data that they need, for example around customer sales volumes, and then access it via the data fabric layer. The key idea is that such requests are satisfied dynamically (even if they require data from several data sources) instead of having to create new data replicas for each new business analytics need.

### What does it do?

Denodo has all the elements needed for such a data fabric architecture and has steadily developed these over its 25 years in business. To actually satisfy data queries dynamically brings with it several challenges. With data scattered among assorted



This **Mutable Quadrant** is derived from 13 high level metrics, the more the image covers a section the better. **Execution** metrics relate to the company, **Technology** to the product, **Creativity** to both technical and business innovation and **Scale** covers the potential business and market impact.

“The Denodo Platform grants the individual wishes of every user, from data scientists to data analysts to data consumers.”  
Carin Velthausz,  
Data Manager,  
ABN-AMRO

systems (at least 400 source systems for a typical large enterprise, according to one survey), there is clearly a mix of source databases and applications. Some data may be tucked away in SAP or Salesforce, other data in specialist marketing or engineering systems, each of which may use different databases and file structures. Denodo has built a kind of uber-optimiser that analyses a query and attempts to satisfy that query in the most appropriate place, along with all the required modern data protection mechanisms. For example, it will try and push computationally intensive queries down to an underlying database, though it also has a dynamic MPP execution engine as a supplement, based on the Presto open-source initiative. Denodo monitors user inquiries and keeps track of which data sources are commonly accessed. All the data source access and the data delivered

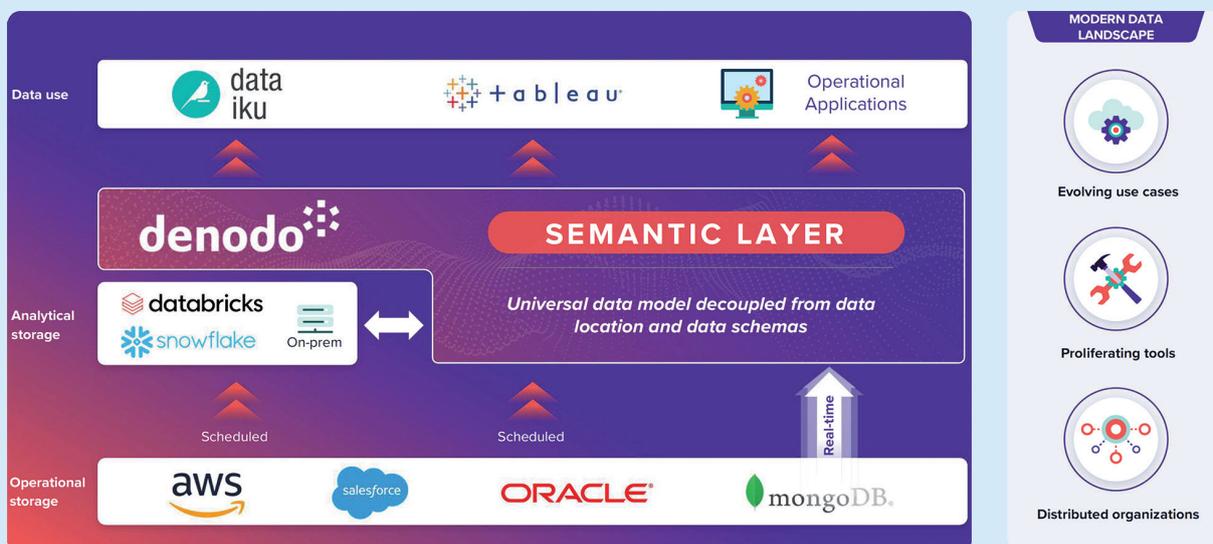


Figure 1 – Logical data architectures embrace the modern data landscape

through Denodo are protected by the embedded policies engine supporting ABAC and RBAC. Its machine-learning approach then suggests appropriate data aggregation and caching to improve performance over time.

This approach does not try to do away with master data management initiatives, where multiple versions of data about customers, products, assets etc are managed and consolidated. Instead, a master data management hub is just another source system as far as Denodo is concerned, and of course, such a source will simplify the task of the Denodo optimiser. End users see a data marketplace based on the underlying catalogue, which itself can apply business policies such as attribute level security, perhaps masking records that contain sensitive data. A workflow engine is also provided, so for example a user might request some data and discover that they need access to it; they can raise a request within the system and this request is assigned to the appropriate security administrator. Once a user has the necessary access then they can invoke a query, either directly within the product or via a third-party tool like Tableau or Qlik. There is a data lineage visual tool that shows the sources of data and for example, whether these have been through some prior transformations.

“Denodo helps us reduce the overall size of the data sets that need to be moved over the network. It also helps us create a logical data fabric that makes our data assets reusable and can be consumed by users spread across the globe.”

Joshua Fletcher,  
Superintendent HSE Reporting  
at BHP

“Data Virtualization has doubled the number of BI projects we completed on time; it reduces cost and enriches BI with new data across internal and external systems.”

Wilson Hung,  
Former IT Director,  
Biogen

### Why should you care?

Denodo also provides a natural language query capability based on ChatGPT, whose large language model is trained on the data definition language that Denodo has access to in its catalogue. The platform monitors queries and provides scheduling capabilities, for example, to set up regular pre-aggregation of certain data that it detects is often accessed. This is a dynamic process, so the system detects new data sources, and can automatically detect changes in underlying database schemas of source systems. Such changes are highlighted in the Denodo system and the catalogue can be adjusted to stay in line with the changed underlying data landscape.

I examined a large number (many dozens) of public customer testimonials of Denodo, and there is no doubt that the technology, which at first glance sounds almost too good to be true, turns out to be well-proven in practice in some very large and challenging customer environments. In this case, the data fabric story does not unravel when you pull on the threads.

### The bottom line

Denodo has been a pioneer of data virtualisation and as such is in a leading position in the emerging data fabric market. It not only has the core components necessary for implementing a data fabric architecture, but these components are mature and proven in many large companies. If your company is considering implementing a data fabric architecture then Denodo is an excellent place to start.

Example

```
SELECT c.state, AVG(s.amount)
FROM customer c JOIN sales s ON c.id = s.customer_id
GROUP BY c.state
```

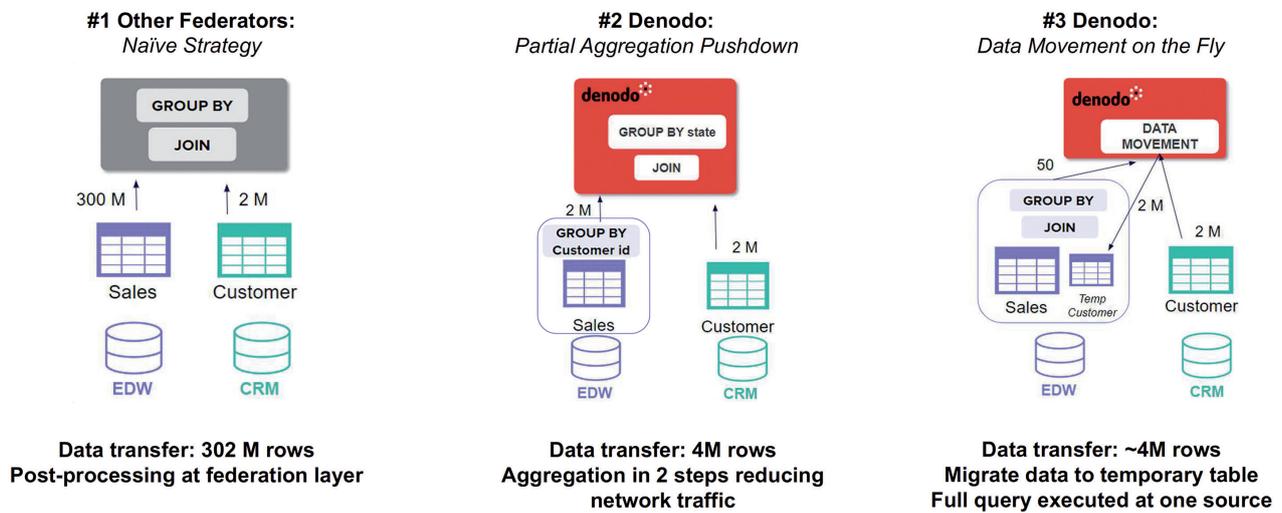


Figure 2 - Real-time execution: multiple sources