DENODO ENTERPRISE PLUS

An advanced logical data fabric platform augmented with AI and semantics to meet complex data integration needs

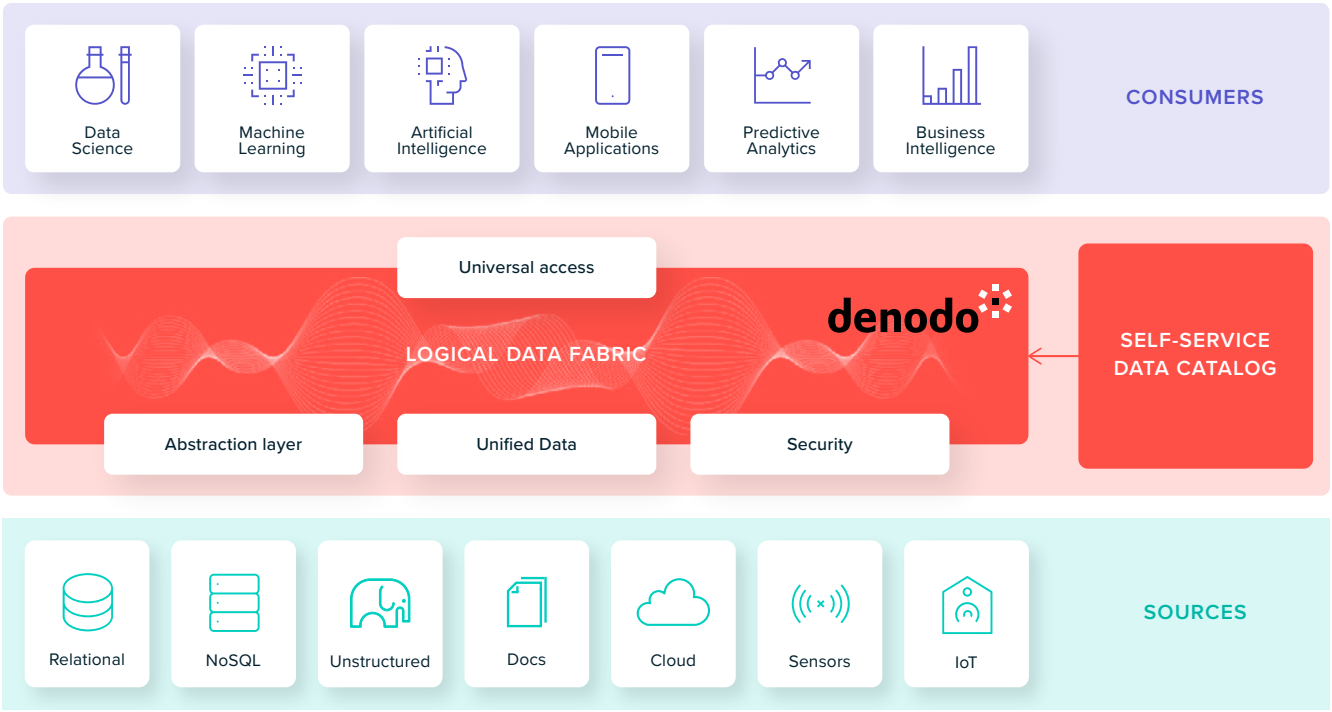

Fig. 1 Logical Data Fabric Architecture

Today's data management landscape is becoming increasingly complex, as data is spread out across many heterogeneous data systems (data warehouses, columnar databases, MPP engines, specialized data stores, SaaS/cloud applications, etc.) that can be in multiple locations (on-premises, cloud, etc.). This makes it difficult to offer a unified view of the data to business applications and to guarantee that governance policies and rules are enforced across the data delivery chain.

Logical data fabric is an architectural style for a unified data delivery platform that abstracts access to multiple data systems for business consumers, hiding the complexity and exposing the data in business friendly formats, while at the same time guaranteeing the delivery of data according to predefined semantics and governance rules, solving today's fundamental data management challenges. A logical data fabric, as defined by leading industry experts TDWI, must have 6 key attributes:

1. Integrate data across multiple cloud environments
2. Automate manual tasks using augmented intelligence
3. Boost analytic performance with rapid data delivery
4. Support data discovery and data science initiatives
5. Provide analyses across data-at-rest and data-in-motion
6. Provide a catalog of all data for discovery, lineage, and associations

Denodo Enterprise Plus provides features to achieve all of the above, and more.

Data virtualization is the key technology for establishing a logical data fabric architecture. The only data integration style designed for distributed architectures, data virtualization provides a logical data access layer on top of multiple heterogeneous systems in hybrid, distributed architectures. Today's data management ecosystem is distributed in nature, so a logical data access layer such as data virtualization is the best fit.

Denodo Enterprise Plus is the only platform that enables organizations to achieve the vision of a logical data fabric. It incorporates features that accelerate the delivery of governed data to business applications, in the most appropriate format for each consumer, across multiple, geographically distributed, heterogeneous systems.

Denodo Enterprise Plus offers strong support for core data virtualization use cases: logical analytical architectures, logical data warehouses, and data services APIs. But it goes beyond traditional data virtualization scenarios to better support new types of users and new types of use cases, such as data science and machine learning (ML) initiatives.

Denodo Enterprise Plus adds mission-critical features, including a recommendations engine, which provides summary recommendations in data discovery scenarios.

Denodo Enterprise Plus also provides AI-powered semantic collaboration features that provide endorsements, warnings, and deprecations, as well as the ability for users to request access, request changes to the metadata, or ask questions about particular data assets.

In addition, Denodo Enterprise Plus represents a big step forward in platform as a service (PaaS) cloud strategy, with capabilities for automatically managing the cloud infrastructure from a centralized web console.

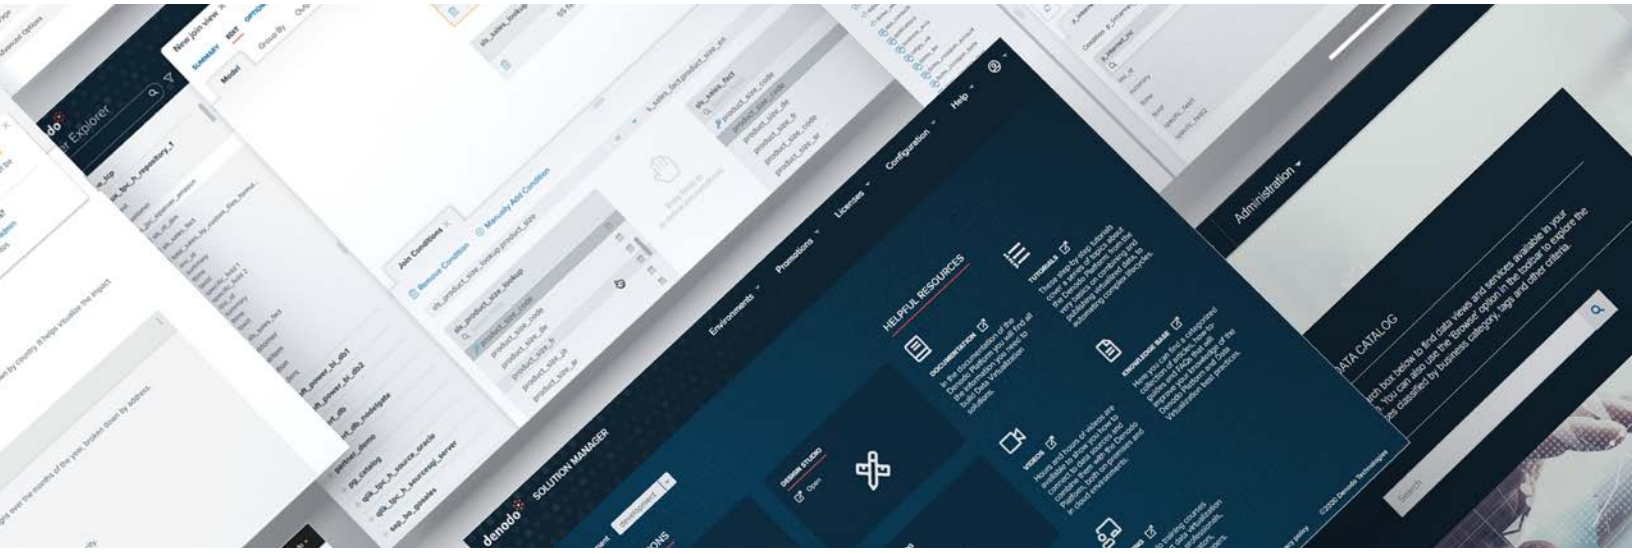

Key Features of Denodo Enterprise Plus

Denodo Enterprise Plus offers an **enhanced, unified user experience with a full web-based interface** for all Denodo Platform components, making it much easier to manage hybrid on-premises/cloud deployments.

The **Denodo Central Web Console**, integrated in the Solution Manager, provides a single entry point to all Denodo tools, for all users, in all Denodo Platform environments, both on-premises and in the cloud. It supports SSO using Kerberos, SAML, OpenID and OAuth, enabling seamless connectivity to all Denodo Platform tools. It also provides tighter integration between tools (e.g., the Diagnostic and Monitoring Tool is now integrated in the Solution Manager).

Developers now benefit from the web-based Design Studio tool with which to develop views and data services (They can still use the desktop version, which is still supported in Denodo Enterprise Plus). This new tool has been designed to maintain and enhance the Denodo Platform's traditional ease-of-use.

In terms of performance, Denodo Enterprise Plus introduces a new concept, Smart Query Acceleration, which accelerates query execution in logical data warehouse/data fabric architecture. Partial aggregates of fact and dimension tables (called "summaries"), which are commonly joined together in many queries of a certain type, are precomputed, and accelerate the execution of future queries. This technique achieves significant performance gains and facilitates the building of ad-hoc queries in a self-service scenario, as these performance enhancements are handled behind the scenes by the Denodo Platform. This extends the Denodo Platform's leadership in distributed query optimization, adding to the strength of other features such as advanced query pushdown, an embedded Presto-based MPP engine, and sophisticated caching mechanisms.

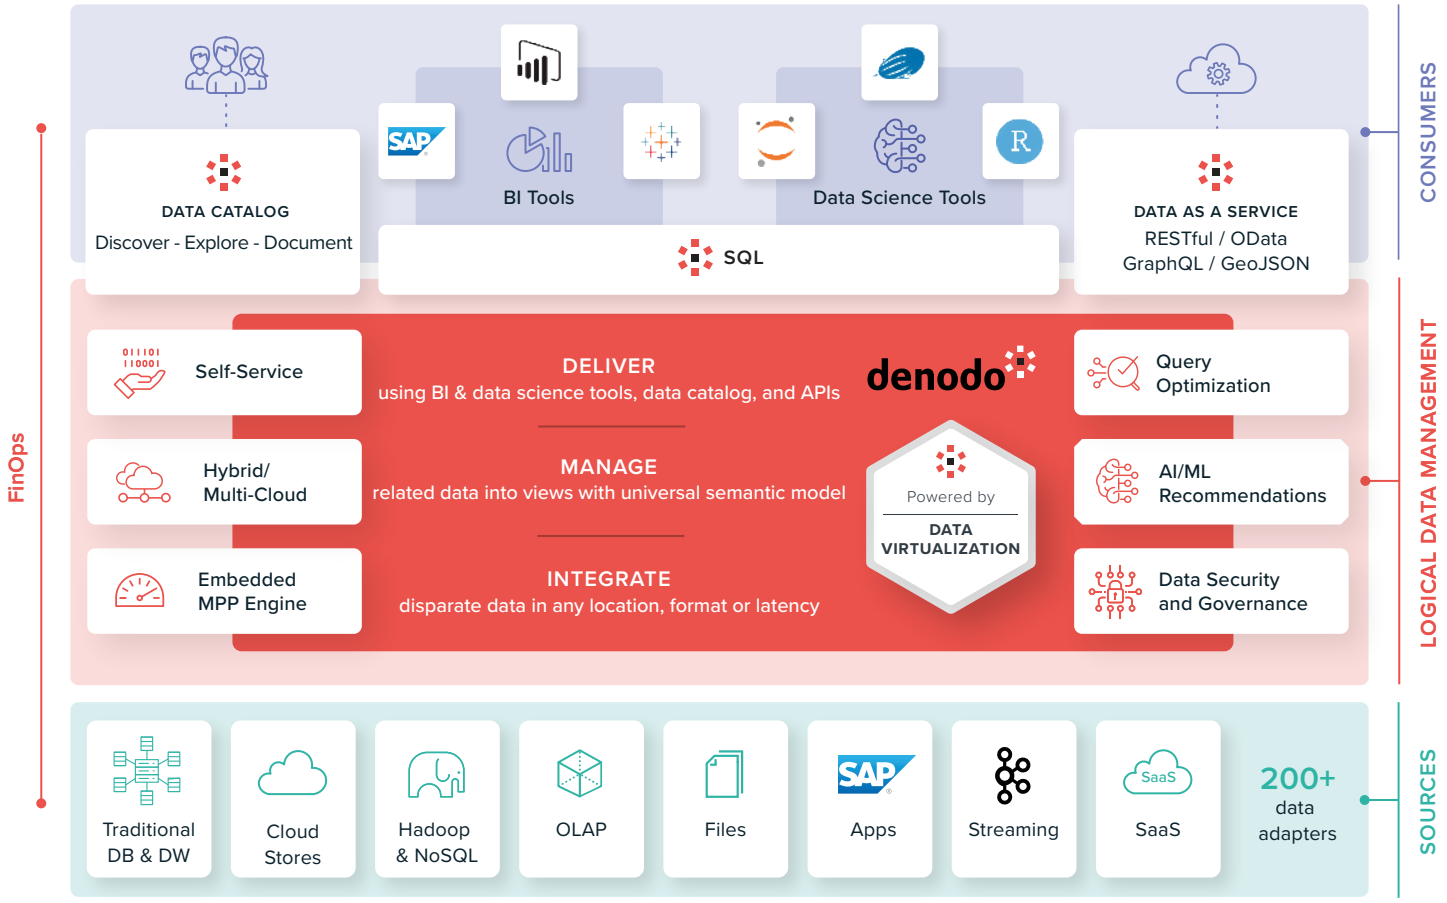

Fig. 2 Denodo Enterprise Plus Architecture

One of the most significant features of Denodo Enterprise Plus is automated infrastructure management in the cloud, which automates all of the tasks related to installing, configuring, deploying, and upgrading Denodo Platform clusters. You can define and configure clusters, including the ability to manage user preferences in such areas as TLS configuration, load-balancing, and autoscaling. Once the clusters have been defined, users can simply press "Start," and the clusters will be automatically installed and created. The Denodo Solution Manager allows for this configuration, as well as integrated monitoring, including **built-in FinOps** for the monitoring and predicting cloud infrastructure costs.

Denodo Enterprise Plus offers advanced support for data services with flexible delivery options (REST, SOAP, OData, and OpenAPI for documentation), the ability to expose data in multiple formats (XML, JSON, HTML, and RSS), and support for the latest security protocols (OAuth, JSON Web Tokens, SAML, Kerberos, HTTPS, HTTP Basic Digest Authentication, or WS-Security). In this area, Denodo Enterprise Plus also adds support for GraphQL, one of the fastest growing data services standards. GraphQL is a query language for APIs. It simplifies the querying of multiple REST endpoints, and with Denodo Enterprise Plus, this is achieved with zero code, and with much better performance, as it leverages the Denodo advanced query optimization engine.

To further facilitate data scientists' access to data through the Denodo Platform, Denodo Enterprise Plus introduces a new tool, the **Denodo Notebook**, which is based on Apache Zeppelin. The Denodo Notebook is fully integrated with the Denodo security system, including SSO support, so all security and data governance policies defined at the Denodo Platform layer are enforced when data scientists use the Denodo Notebook.

For business users, the **Denodo Data Catalog** includes new features such as automatic recommendations; enhanced collaboration with endorsement, comments, etc.; and enhanced profiling and search features (smart ranking). You can also use the Denodo Catalog to create exports to CSV, Excel, and Tableau Data Extracts. The catalog also shows usage statistics, so you know who used what data sets when, and how they used the data.

KEY FEATURES OF DENODO ENTERPRISE PLUS

A full web-based interface for all Denodo tools with SSO support: an integrated, web-based experience across all tools.

Web-based Design Studio tool for developers, providing ease-of-use across all of the steps in the data-service development process.

Smart query acceleration: optimizes query execution through intelligent caching, cost-based optimization, and an AI-recommendations engine.

Embedded Presto MPP Engine: provides data consumers with a highly scalable and performant SQL engine to access data lake data while simplifying data access through an intuitive GUI interface.

Automated infrastructure management for the cloud: PaaS support, including cluster configuration (TLS, load-balancing, autoscaling, etc.), start/stop controls, automatic installation of updates, and integrated monitoring.

GraphQL support: zero-code creation of GraphQL data APIs, with first-class performance leveraging the Denodo query optimizer.

Attribute-based accesscontrol (ABAC) combines the semantic layer, tagging, and session attributes to allow highly granular access policies, defined globally based on various factors including data sensitivity, user roles, and business requirements.

Data Catalog with automatic recommendations and enhanced collaboration, profiling, and search features.

- **AI driven automatic recommendation of datasets:** This new capability enables the Data Catalog to display personalized recommendations to users, based on their previous activity.
- **Collaboration using the Data Catalog:** Users can add endorsements, warnings, and deprecation notes about views and web services. This way, data stewards can better communicate with business users.

Business Benefits

HIGHLY ECONOMICAL
Integrate data easily and reliably, in real time, at a fraction of the cost of traditional data integration approaches such as extract, transform, and load (ETL) processes, and with significantly greater agility.

UNDERSTAND AND MANAGE CLOUD COST
Manage and predict future cloud infrastructure costs and understand how to optimize cost and performance by leveraging an interactive dashboard highlighting active metadata on cloud infrastructure and egress charges.

FASTER PATH TO VALUE
Deliver contextual, reliable information, faster, for more actionable insights. Agile enterprises are industry leaders.

LOCATION-AGNOSTIC:
Deploy in any location – on-premises, cloud, or edge – without sacrificing scale or data governance capabilities.

BUSINESS-FRIENDLY
Abstract the complexity of modern data ecosystems (myriad sources, multiple formats, distributed, heterogeneous, diverse) from business. Expose data in the right format and use the naming conventions required by every type of user and application at almost no cost. Rapidly adjust to changes in requirements.

EXTENDED COLLABORATION OPTION
Foster collaboration and ensure business stakeholders are maximizing the potential of the platform with a built-in feedback loop, allowing users to request changes, report data quality issues, request access, or ask a question on specific data products.

ENTERPRISE GRADE
Support multiple lines of business, multiple use cases, multiple personas, and thousands of users.

A DIGITAL MARKETPLACE:
Enable a digital marketplace that empowers a community of analysts to find and use information assets quickly, which is essential in this age of self-serve analytics.

BUILT-IN INTELLIGENCE: EXTEND THE POWER OF THE PLATFORM

Semantic powered AI: collaboration feature providing endorsements, warnings, and deprecation, significantly improves data governance by proactively informing the user of changes/needs.

AI based performance recommendations: help reduce the burden on IT infrastructure by optimizing the computation and execution of queries.

AI based recommendations and summaries: enhance the experience for data scientists and analysts, enabling a higher level of self-service and sufficiency, significantly improving time-to-discovery while also reducing the data engineering cost.

COMPARING THE VARIOUS SUBSCRIPTION TIERS

There are several flexible subscription options to choose from, designed to suit the needs of various projects, from small departmental projects to enterprise-wide digital transformation.

	DENODO PROFESSIONAL Small, single-use-case projects within individual departments	DENODO STANDARD Multiple use cases within individual departments	DENODO ENTERPRISE Enterprise-wide deployment for multiple use cases and groups and large data volumes	DENODO ENTERPRISE PLUS Comprehensive automation, collaboration, and advanced security for enterprise-wide deployments
Number of Data Sources Supported	5	Unlimited	Unlimited	Unlimited
FinOps Logging and Integrated Dashboard	✓	✓	✓	✓
Integration with Version Control Systems (VCS)		✓	✓	✓
Smart Query Acceleration using summaries			✓	✓
VQL procedures			✓	✓
Advanced Diagnostic & Monitoring Tool			✓	✓
Data Catalog			✓	✓
Integration with external Massive Parallel Processing (MPP) engines like Impala, Spark, and others			✓	✓
Integration with 3rd party data modeling tools (ER/Studio, Erwin, etc.)			✓	✓
Embedded Presto-based MPP engine				✓
Automatic recommendation of summaries				✓
Global security policies				✓
Import data governance tags from external catalogs				✓
Data Catalog: Natural language queries and AI-based recommendations				✓
Data Catalog: Dataset collaboration through endorsements, warnings, and deprecation notes				✓

Denodo Enterprise Plus Capabilities Sheet

DATA SOURCES

Relational Databases
- Generic (JDBC)
- IBM DB2 (JDBC): 8, 9, 10, 11, 12 for LUW, 9,10 for z/OS, AS400
- Multi Layered Denodo deployments (JDBC): 5.5, 6.0, 7.0, 8.0
- Apache Derby (JDBC): 10
- Informix (JDBC): 7, 12
- MS SQL*Server (JDBC, ODBC): 2000, 2005, 2008, 2008R2, 2012, 2014, 2016, 2017
- MySQL (JDBC): 4, 5, 8
- Oracle (JDBC): 8i, 9i, 10g, 11g, 12c, 18c, 19c
- Oracle E-Business Suite (JDBC): 12
- PostgreSQL (JDBC): 8, 9, 10, 11, 12
- Sybase Adaptive Server Enterprise (JDBC): 12, 15
- MS Access (ODBC)

In-Memory Databases
- SAP HANA (JDBC): 1,2
- Oracle TimesTen (JDBC): 11g
- Oracle 12c In-Memory

Parallel Databases and Appliances
- Exasol (JDBC)
- GreenPlum (JDBC): 4.2
- HP Vertica (JDBC): 7, 9
- Netezza (JDBC): 4.6, 5.0, 6.0, 7.0
- Oracle Exadata (JDBC): X5-2
- ParAccel 8.0.2 (by using ParAccel 2.5.0.0 JDBC3g with SSL driver)
- SybaseIQ (JDBC) 12.x, 15.x
- Teradata (JDBC): 12, 13, 14, 15, 16, 17
- Yellowbrick

Cloud Data Warehouse / RDBMS
- Alibaba ApsaraDB for OceanBase MySQL (JDBC)
- Alibaba ApsaraDB for OceanBase Oracle (JDBC)
- Alibaba ApsaraDB RDS for MySQL (JDBC)
- Alibaba ApsaraDB RDS for PostgreSQL (JDBC)
- Alibaba ApsaraDB RDS for Microsoft SQL Server (JDBC)
- Alibaba ApsaraDB PolarDB for MySQL (JDBC)
- Alibaba ApsaraDB PolarDB for PostgreSQL (JDBC)
- Alibaba ApsaraDB AnalyticDB for MySQL (JDBC)
- Alibaba ApsaraDB AnalyticDB for PostgreSQL (JDBC)
- Amazon Redshift (JDBC)
- Amazon Athena (JDBC)
- Amazon Aurora (JDBC)
- Amazon DynamoDB
- Azure Cosmos DB
- Azure SQL Database
- Azure Synapse Analytics
- Delta Lake
- Google AlloyDB (JDBC)
- Google Big Query (JDBC)
- Google Cloud SQL (JDBC)
- Google Spanner DatBoost (JDBC)
- MongoDB Atlas
- Snowflake (JDBC)

Big Data
- Apache Hive (JDBC): 0.12, 1.1.0, 1.1.0 for Cloudera, 1.2.1 for Hortonworks, 2.0.0
- Impala (JDBC): 2.3
- Spark SQL (JDBC): 1.5, 1.6, 2.x, 3.x
- PrestoDB (JDBC)
- PrestoSQL / Trino (JDBC)
- Databricks Delta 2.x

NoSQL
- MongoDB
- Cassandra
- HBase

Multi-Dimensional Sources
- SAP BW (BAPI/XMLA): 3.x
- SAP BI 7.x (BAPI): 7.x
- Mondrian (XMLA): 3.x
- IBM Cognos TM1
- MS SQL Server Analysis Services 200x
- Essbase (XMLA): 9, 111

Data Lake Storage
- S3
- Azure Data Lake Storage
- Azure Data Lake Storage Gen 2
- Azure Blob Storage
- Google Cloud Storage
- Parquet
- Avro

Web Services
- SOAP
- REST (XML, RSS, ATOM, JSON)
- OData

Flat and Binary Files
- CSV, pipe-delimited, regular expression-parsed
- MS Excel xls 97-2003
- MS Excel xlsx 2007 or later
- MS Access
- XML
- JSON
- SAS Files (SAS7BDAT)
- All files can be local or in remote filesystems, through FTP/ SFTP/FTPS, and in clear, zipped and/or encrypted format.

Indexes and unstructured content
- CMS, file systems, text
- ElasticSearch 6.4, 6.7

Cloud, SaaS, Web Sources with Simplified OAuth Security
- Adobe Analytics
- Amazon

- Google
- Google Sheets
- Facebook
- LinkedIn
- MS Azure Data Lake
- MS Sharepoint (via OData)
- MS Dynamics 365 Business Central / Customer Engagement
- Marketo
- ServiceNow
- Salesforce (SOQL)
- Twitter
- Workday
- many more through configurable JSON and XML adapters

Active Directory as Source or Leveraging Security
- LDAP v3
- Microsoft Active Directory 2003, 2008

Message Queues
- MQSeries
- SonicMQ
- ActiveMQ
- Tibco EMS
- Other JMS compatible services

Semantic Repositories
- Semantic repositories in Triple Stores / RDF accessed through SPARQL endpoints.

Packaged Applications
- SAP ERP/ECC (BAPIS and tables)
- Oracle E-Business Suite 12
- Siebel
- SAS (SAS JDBC Driver): 7 and higher

Mainframe
- IMS
  - IBM IMS native drivers: 8, 9
  - IMS Universal Drivers: 11

Hierarchical Databases
- Adabas (SOA Gateway and Denodo's SOAP connector): 5, 6

Denodo SDK for Custom Connectors

PUBLISHING OPTIONS

- SQL Based access via JDBC, ODBC and ADO.NET
- Web Services
  - REST
  - OData
  - Open API (a.k.a Swagger)
  - GraphQL
  - SOAP
- OAuth, OAuth 2.0 (JWT)
- SAML
- SSL
- WS-Security
- JMS listeners for message queues
- Denodo Scheduler for batch process and lite ETL

DATA CATALOG

- Web-UI for data discovery and exploration for business users
- Business Categories and Business Tags
- Full search capabilities on metadata and actual data
- Integrated request management (access, changes, data quality issues, etc.)
- Query wizards for customizing datasets
- Full-featured SQL shell allows execution of complex queries from the Data Catalog web-UI
- Export to CSV, Excel, and Tableau Data Extracts
- Query sharing features
- Self-service dataset export to a shared sandbox for IT review before final publication for global use
- Automatic personalized recommendations and shortcuts to recommended datasets
- Endorsement of datasets, comments, warnings, etc.
- Usage statistics: who uses what data, when and how
- Profiling information
- Intelligent search with smart ranking of results

PERFORMANCE OPTIMIZATIONS

- Smart Query Acceleration for Analytics
  - Aggregate Aware Summaries
- Massive Parallel Processing (MPP) integration for Query Acceleration and Caching
- Full and partial aggregation and join pushdown, even in federated views
- Support for alternative data sources
- On-the-fly data movement for optimization
- Cost Based Optimization (data statistics, data source indexes, data source execution model and parameters, network transfer rates)
- Pushdown of selections/projections/joins/groupby operations also on federated views
- Multiple join strategies
- Simplifying partitioned unions (Partition pruning)
- and many more

CACHE AND DATA MOVEMENT OPTIONS

- Multi-mode caching: full, partial, incremental, or total refresh, event-based or scheduled, configured at the view level, incremental queries for SaaS sources
- Amazon Athena
- Azure SQL
- Azure SQL Data Warehouse
- Azure Synapse Analytics
- Amazon Redshift
- Databricks 2.x
- Delta Lake
- IBM DB2 (8, 9, 10, 11 for LUW, 9,10,11 for z/OS)
- Hive 2.0.0
- Impala
- MS SQL Server (2000, 2005, 2008, 2008R2, 2012, 2014, 2016, 2017)
- MySQL (4 and 5)
- Netezza (6 and 7)

- Oracle (8i, 9i, 10g, 11g, 12c, 12c in-memory, 18c, 19c)
- Oracle TimesTen 11g
- PostgreSQL (9 and 10)
- Presto
- SAP HANA
- Snowflake
- Spark (1.5,1.6 and 2.x)
- Teradata (12, 13, 14, 15, and 16)
- Vertica (7 and 9)
- Yellowbrick
- Configurable “generic” adapter for other databases with JDBC drivers

DATA PIPELINES

- Remote Tables (created through UI or stored procedure)
- Denodo Scheduler

EMBEDDED MPP

- MPP engine based on Presto to accelerate access to data lake
- Graphical introspection of object storage (S3, ADFS, GFS, HDFS, etc.)
- Support for Parquet, Delta, and Iceberg
- New advanced optimization techniques to federate data lake content with any other data source

THIRD-PARTY MPP OPTIONS

- Impala
- Presto
- Spark 1.5, 1.6, 2.x
- Databricks 2.x

DATA GOVERNANCE

- Data source refresh, change impact analysis, dependency tree, full data lineage
- Denodo Governance Bridge: integration with IBM Information Governance Catalog
- API to publish metadata and lineage information to data governance tools like Informatica EDC, Collibra, etc.

SECURITY

Data in Motion – secure channels
- Using SSL/TLS
- Client-to-Denodo and Denodo-to-source
- Available for all protocols (JDBC, ODBC, ADO.NET and WS)

Data at Rest - secure storage
- Cache: third party database. Can leverage its own encryption mechanism
- Swapping to disk: serialized temporarily stored in a configurable folder that can be encrypted by the OS

Encryption/Decryption
- Support for custom decryption for files and web services
- Transparent integration with RDBMs encryption
- Encrypted metadata import/export

User and Role Based including integration with AD/ LDAP
- Row and Column level authorization
- Advanced customizable masking
- Custom policies for specific security constraints and integration with external policy servers

Global Policies
- Tag-based security policies
- Support for RBAC and ABAC
- Column and row restrictions, multiple masking options, deny execution

Authentication
- Native and LDAP/Active Directory based Support for Kerberos and Windows SSO
- Base64
- Kerberos
- NTLM
- OAuth, OAuth 2.0 (JWT)
- SAML
- Two-factor authentication (through supported identity providers: Okta, Duo, etc.)
- SSL
- WS-Security
- Pass-through session credentials to leverage existing source privileges

DATA MODELING

- Design Studio: Web-based development studio for data modeling
- Desktop version also available
- Bottom-Up and Top-Down (through Interface Views)
- Integration with third-party modeling tool
  - ER/Studio Data Architect
  - ERwin Data Modeler
  - IBM InfoSphere Data Architect
  - SAP PowerDesigner

DATA QUALITY

- Library of transformation, filter and matching functions and quality rules for validating, cleansing, enriching, standardizing, matching and merging data
- Extensible through Custom Functions
- Integration with external DQ tools

MONITORING

- Denodo Diagnostic and Monitoring Tool (DMT) integrated in the Solution Manager
- Extensible usage and metadata dashboards integrated in Apache Superset
- FinOps dashboard to monitor and understand key metrics associated with cloud costs like egress,

query cost, etc.
- Detailed monitoring information is available in logs for integration with log management tools like Splunk, ELK, Cloudwatch, etc.
- AI-based recommendations
- Performance recommendations for Smart Query acceleration using usage history
- Personalized data discovery recommendations in the Data Catalog
- Monitoring is also available via SNMP and JMX standards. Therefore interoperate with most leading Systems Management packages (e.g., HP OpenView, Nagios, Zenoss, Osmius, IBM Tivoli and Microsoft WinRM)

OPERATIONS

- Solutions Manager to automate operations and promotions tasks
  - Centralized management and distribution of updates to clients
  - Centralized management of license keys
  - Define promotion revisions and their dependencies and deploy them to a production cluster with zero downtime
  - Centralized management of data source properties and logs
  - REST API for automation of tasks from DevOps tools (e.g. jenkins)
- Integrated Infrastructure Management for Cloud (AWS)
  - Creation and management of clusters: define type of EC2 instances, number of EC2 instances, etc.
  - Creation of load balancers and Auto Scaling groups.
  - Installation and launch of the Denodo servers.
  - Update of Denodo version
  - Enable SSL in the Denodo servers.
- Multi-User Development with Version Control integration
  - Subversion
  - Microsoft TFS
  - Git
- Resource Manager to limit and allocate resources to each session, role or user in a way that optimizes resources utilization for each application
  - Change resources priority
  - Enforce limited timeouts or limits on number of rows
  - Add daily quotas per minute/day/month: e.g. only 50 queries per day

DEPLOYMENT PATTERNS

- On-premises, private cloud, public cloud
  - On-premises, private cloud, public cloud
  - Basic single server configuration
  - HA cluster with load balancing (Active-Passive and Active-Active)
  - Shared or distributed local cache
  - Geographically distributed server environments
  - Multiple Denodo instances peer-to-peer or multi-layered
  - Containerization support through Docker
- Public cloud
  - Denodo Platform for AWS
  - Denodo Platform for Azure
  - Denodo Platform for GCP
- Auto-scaling support both in AWS and Azure

USER INTERFACES

- Central Web Console: integrated access to all Denodo UIs with SSO (Kerberos, SAML, OpenID and OAuth)
- Solution Manager: Centralized UI for administrators to manage deployments and promotions, including automatic management of cloud infrastructure (AWS)
- Design Studio: Web-based Development Studio, drag-and-drop and low-code developer studio geared to data-oriented developers such as data engineers, power users, and citizen integrators; publish data services with a few clicks.
- Desktop Dev. Studio (VDP Admin tool)
- Data Catalog: Easy-to-use web-based interface for business-oriented users such as data stewards, data analysts, and citizen analysts.
- Diagnostic and Monitoring Tool: centralized Web-based UI for monitoring, auditing, and troubleshooting for data engineers and administrators.

OPERATING SYSTEMS

- Microsoft Windows (32-bit and 64-bit platforms): Windows Server 2019, Windows Server 2016, Windows Server 2012, Windows Server 2008, Windows 10, Windows 8.1, and Windows 7
- Linux (32-bit and 64-bit platforms): Red Hat Enterprise Linux (RHEL) 6 and 7, Oracle Linux 6 and 7, Ubuntu 12.04 LTS and 14.04 LTS, CentOS 6 and 7
- Any Java 11 or greater compatible OS

MINIMUM HARDWARE REQUIREMENTS

- Processor: Intel Xeon quad-core or similar. High-load scenarios or cases with complex calculations may require 8 cores or more.
- Physical memory (RAM): 16 gigabytes of memory so the Denodo server can allocate a runtime heap space up to 8 gigabytes.
- Disk space: Minimum: 5 gigabytes, Recommended: 100 gigabytes. Denodo only needs around 1 GB of disk space. If the cache is installed on the same server, more disk space will be required..