

# Denodo Standard 8.0 Datasheet



The accelerating adoption of cloud platforms, cloud-based data stores, and software-as-a-service (SaaS) applications introduces new complexities and challenges for organizations that need to integrate their cloud data with on-premises databases and applications so that all systems can work together seamlessly. The accelerating complexity of data in today's hybrid, multi-cloud world means that organizations need a way to connect everything together in a way that is flexible and agile. However, it can be challenging and resource intensive to integrate and make all of the data available in real time for analytics or data services. Traditional approaches, which use extract, transform, and load (ETL) processes to copy all of the data into a new repository for analysis, are time consuming and expensive, and they can result in data that is fragmented and not up-to-date.

Through data virtualization, Denodo offers a better solution that integrates data in real time, directly from the data sources, without having to move the data to a new, consolidated location. Denodo's no-code/low-code, modern data virtualization solution provides a data abstraction layer that accelerates access to data, while providing security and governance. The Denodo Platform can be used to speed analytics by enabling the creation of a logical data warehouse, which can seamlessly combine data from cloud and on-premises warehouses, in real time. The data from all of the connected sources can also be made available as data services, through the Denodo Platform's no-code data API creation capabilities. Denodo integrates and delivers data in as little as the time with as much as 80% cost savings compared with other approaches.

Denodo Standard 8.0 brings the power of Denodo's modern data virtualization to project-specific uses by providing an easier and more cost effective starting point. Available pay-as-you-go in leading cloud marketplaces with automated setup and configuration, it lowers the cost and accelerates data access for analytics and data services scenarios. This makes it a great fit for organizations of any size that are looking for ROI without making a large financial commitment.


# Key Features of Denodo Standard 8.0

Denodo Standard 8.0 is built on the same proven high-performance Denodo data virtualization technology used in Denodo Platform 8.0 to quickly connect to data and make it available for analysts, data scientists, and consuming applications. Denodo Standard 8.0 comes with an easy-to-use unified Web interface and automated cloud infrastructure setup and management to make it fast and efficient for small teams to quickly get into production.

Denodo Standard 8.0 offers all the capabilities necessary to get data integration projects into production with minimal effort and resources, within the budgets of small teams or initial projects. With out-of-the-box connectivity to more than 150 data source types, a visual design studio to combine, integrate and normalize data and smart query acceleration, Denodo Standard 8.0 enables you to connect, combine and consume data quickly and efficiently.

Developers benefit from the productivity offered by the new Web-based Design Studio tool, which makes it possible to quickly develop views and data services without writing any code. The new Denodo Design Studio enhances Denodo Standard's ease-of-use and supports SSO to make it securely available to your team members, wherever they may be. The Design Studio enables more of your team to contribute through the creation of data views using no-code data modeling. It supports single sign-on (SSO) using Kerberos, SAML, OpenID and OAuth, enabling seamless connectivity to all Denodo tools.

Denodo Standard 8.0 offers advanced support for data services with flexible delivery options (REST, SOAP, OData and OpenAPI), the ability to expose data in multiple formats (XML, JSON, HTML, RSS) and unified support for the latest security protocols (OAuth, JSON Web Tokens, SAML, Kerberos, HTTPS, HTTP Basic Digest Authentication, or WS-Security). Denodo 8.0 also adds support for the popular GraphQL, to simplify the queries of multiple REST endpoints.



Denodo Standard 8.0 benefits from Denodo’s performance leadership by utilizing powerful capabilities including distributed query optimization, advanced query pushdown, and sophisticated caching capabilities.

Denodo Standard 8.0 includes automated cloud infrastructure management to accelerate all of the tasks related to installing, configuring, deploying, and upgrading Denodo Platform clusters.

Denodo Standard 8.0 includes the Apache Zeppelin-based Denodo Notebook for data scientists, leveraging security and SSO. This enables data scientists to more quickly find the right data and put it to work in Machine Learning and predictive analytics models.




Fig. 1 Denodo Standard 8.0 Architecture

## KEY FEATURES OF DENODO STANDARD 8.0 INCLUDE:



**Unified, Web-based interface to all Denodo Standard tools**, with SSO support.



**150+ out-of-the-box connectors** to quickly connect sources and targets in real time.



**Web-based Design Studio** for developer ease-of-use in creating data views by combining data and developing data services. Provides visual data lineage to understand data sources and transformations.



**High-performance query acceleration** of virtualized data views to speed responsiveness while minimizing the loads on source systems and network traffic.



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



**Advanced Data Services:** Zero-code, automated creation of data APIs, including the ability to use GraphQL for efficient data services.



**Apache Zeppelin-based Denodo Notebook:** Integrated with Denodo SSO to better enable data scientists with data from connected systems.

## Business Benefits

### FAST

Get started in minutes by leveraging automated cloud infrastructure setup and a highly productive no-code/low-code tool with easy connectivity to over 150 data sources.

### FLEXIBLE, PAY-AS-YOU-GO PRICING

Get started with a minimal monetary commitment that only requires payment when you use Denodo Standard.

### BUSINESS-FRIENDLY


Denodo Standard is an easy-to-use solution that abstracts the complexity of modern data ecosystems (distributed, heterogeneous, diverse data sources) from business users by making all the connected data appear to be a single database.

### EFFICIENT, REAL-TIME DATA ACCESS

Efficient, real-time data access: Integrate cloud and on-premises data in a fraction of the time and cost of traditional data integration approaches such as ETL processes.

### FASTER PATH TO VALUE


The low initial investment on on-going charges delivers rapid ROI by delivering data faster and more cost effectively. This makes Denodo Standard a great fit for small and medium sized businesses just starting out with data virtualization.



## High Level Features of Denodo Standard

- ✓ Denodo data virtualization engine
- ✓ Hourly and annual pricing available via cloud marketplace subscription.
  - ✓ Production instance comprises of Denodo's single server licensing
  - ✓ Solution manager is included with the subscription
  - ✓ Single development instance is included in the pricing while a UAT instance is priced separately.
- ✓ Support for up to 5 data sources
- ✓ Comprehensive security support (authentication, authorization, user and role based policy, masking, encryption and more)
- ✓ Denodo notebook, data lineage capabilities are included in the subscription.
- ✓ Core performance capabilities such as standard query optimizations are all included.

Users looking for broader data management capabilities beyond data virtualization can leverage the Denodo Platform to gain access to features such as unlimited data sources, data catalog, smart query acceleration (performance summaries), advanced diagnostics, and version control.



# Denodo Standard 8.0 Capabilities

<p><b>DATA SOURCES</b></p> <p><b>Relational Databases</b></p> <ul style="list-style-type: none"> <li>• Generic (JDBC)</li> <li>• IBM DB2 (JDBC): 8, 9, 10, 11, 12 for LUW, 9,10 for z/OS, AS400</li> <li>• Multi Layered Denodo deployments (JDBC): 5.5, 6.0, 7.0, 8.0</li> <li>• Apache Derby (JDBC): 10</li> <li>• Informix (JDBC): 7, 12</li> <li>• MS SQL Server (JDBC, ODBC): 2000, 2005, 2008, 2008R2, 2012, 2014, 2016, 2017</li> <li>• MySQL (JDBC): 4, 5</li> <li>• Oracle (JDBC): 8i, 9i, 10g, 11g, 12c, 18c, 19c</li> <li>• Oracle E-Business Suite (JDBC): 12</li> <li>• PostgreSQL (JDBC): 8, 9, 10, 11</li> </ul> <p><b>Sybase Adaptive Server Enterprise (JDBC): 12, 15</b></p> <ul style="list-style-type: none"> <li>• MS Access (ODBC) In-Memory Databases</li> <li>• SAP HANA (JDBC): 1</li> <li>• Oracle TimesTen (JDBC): 11g</li> <li>• Oracle 12c</li> </ul> <p><b>In-Memory Parallel databases and appliances</b></p> <ul style="list-style-type: none"> <li>• Greenplum (JDBC): 4.2</li> <li>• HP Vertica (JDBC): 7, 8</li> <li>• Oracle Exadata (JDBC): X5-2</li> <li>• ParAccel 8.0.2 (by using ParAccel 2.5.0.0 JDBC3g with SSL driver)</li> <li>• Netezza (JDBC): 4.6, 5.0, 6.0, 7.0</li> <li>• SybaseIQ (JDBC) 12.x, 15.x</li> <li>• Teradata (JDBC): 12, 13, 14, 15</li> <li>• Yellowbrick</li> </ul> <p><b>Cloud Data Warehouse / RDBMS</b></p> <ul style="list-style-type: none"> <li>• Amazon Redshift (JDBC)</li> <li>• Amazon Athena (JDBC)</li> <li>• Amazon Aurora (JDBC)</li> <li>• Amazon DynamoDB</li> <li>• Amazon RDS</li> <li>• Azure Cosmos DB</li> <li>• Azure SQL Database</li> <li>• Azure Synapse Analytics</li> <li>• Delta Lake</li> <li>• Google Cloud SQL</li> <li>• Google BigQuery (JDBC)</li> <li>• MongoDB Atlas</li> <li>• Snowflake (JDBC)</li> </ul>	<p><b>Big Data</b></p> <ul style="list-style-type: none"> <li>• Apache Hive (JDBC): 0.12, 1.1.0, 1.1.0 for Cloudera, 1.2.1 for Hortonworks, 2.0.0</li> <li>• Impala (JDBC): 2.3</li> <li>• Spark SQL (JDBC): 1.5, 1.6</li> <li>• Presto (JDBC)</li> <li>• Databricks 2.x</li> </ul> <p><b>NoSQL</b></p> <ul style="list-style-type: none"> <li>• MongoDB</li> <li>• Cassandra</li> </ul> <p><b>Multi-Dimensional Sources</b></p> <ul style="list-style-type: none"> <li>• SAP BW (BAPI/XMLA): 3.x</li> <li>• SAP BI 7.x (BAPI): 7.x</li> <li>• Mondrian (XMLA): 3.x</li> <li>• IBM Cognos TM1</li> <li>• MS SQL Server Analysis Services 200x</li> <li>• Essbase (XMLA): 9, 11</li> </ul> <p><b>Data Lake Storage</b></p> <ul style="list-style-type: none"> <li>• S3</li> <li>• Azure Data Lake Storage</li> <li>• Azure Data Lake Storage Gen 2</li> <li>• Azure Blob Storage</li> <li>• Google Cloud Storage</li> <li>• Parquet (Distributed File System Connector)</li> <li>• Avro</li> </ul> <p><b>Web Services</b></p> <ul style="list-style-type: none"> <li>• SOAP</li> <li>• REST (XML, RSS, ATOM, JSON)</li> </ul> <p><b>Flat and Binary Files</b></p> <ul style="list-style-type: none"> <li>• CSV, pipe-delimited, Regular expression-parsed</li> <li>• MS Excel xls 97-2003</li> <li>• MS Excel xlsx 2007 or later</li> <li>• MS Access</li> <li>• XML</li> <li>• JSON</li> <li>• SAS Files (SAS7BDAT)</li> <li>• All files can be locally accessible or in remote filesystems, through FTP/ SFTP/FTPS, and in clear, zipped and/or encrypted format.</li> </ul> <p><b>Indexes and unstructured content</b></p> <ul style="list-style-type: none"> <li>• CMS, file systems, pdf, MS Word, text, email servers, knowledge bases, indexes</li> <li>• Elastic Search 6.4, 6.7</li> </ul>	<p><b>Cloud, SaaS, Web Sources with Simplified OAuth Security</b></p> <ul style="list-style-type: none"> <li>• Amazon</li> <li>• Google</li> <li>• Google Sheets</li> <li>• Facebook</li> <li>• LinkedIn</li> <li>• MS Azure Data Lake</li> <li>• MS SharePoint (by using the OData connector)</li> <li>• MS Dynamics 365 Business Central / Customer Engagement</li> <li>• Marketo</li> <li>• ServiceNow</li> <li>• Salesforce (SOQL)</li> <li>• Twitter via APIs with simplified OAuth integration (1.0, 1.0a and 2.0)</li> <li>• Workday</li> </ul> <p><b>Active Directory as source or leveraging security</b></p> <ul style="list-style-type: none"> <li>• LDAP v3</li> <li>• Microsoft Active Directory 2003, 2008</li> </ul> <p><b>MS Queues as data source and Delivery</b></p> <ul style="list-style-type: none"> <li>• MQSeries</li> <li>• SonicMQ</li> <li>• ActiveMQ</li> <li>• Tibco EMS</li> </ul> <p><b>Semantic Repositories</b></p> <ul style="list-style-type: none"> <li>• Semantic repositories in Triple Stores / RDF accessed through SPARQL endpoints.</li> </ul> <p><b>Packaged Applications</b></p> <ul style="list-style-type: none"> <li>• SAP ERP/ECC (BAPIS and tables)</li> <li>• Oracle E-Business Suite 12</li> <li>• Siebel</li> <li>• SAS (SAS JDBC Driver): 7 and higher</li> </ul> <p><b>Denodo SDK for Custom Connectors</b></p> <ul style="list-style-type: none"> <li>• CouchDB</li> <li>• Lotus Domino</li> </ul> <p><b>Mainframe</b></p> <ul style="list-style-type: none"> <li>• IMS <ul style="list-style-type: none"> <li>• IBM IMS native drivers: 8, 9</li> <li>• IMS Universal Drivers: 11</li> </ul> </li> </ul>
---	--	--

<p><b>Hierarchical databases</b></p> <ul style="list-style-type: none"> <li>• Adabas(SOA Gateway and Denodo's SOAP connector): 5, 6</li> </ul> <p><b>Additional data sources:</b></p> <ul style="list-style-type: none"> <li>• Apache Solr</li> <li>• Kafka Messages</li> <li>• Hadoop HBase</li> <li>• Hadoop HCatalog</li> <li>• Hadoop HDFS</li> <li>• IBM BigInsights</li> <li>• Pivotal HAWQ</li> </ul> <p><b>PUBLISHING OPTIONS</b></p> <ul style="list-style-type: none"> <li>• SQL Based access via JDBC, ODBC and ADO.NET</li> <li>• Web Services <ul style="list-style-type: none"> <li>• SOAP</li> <li>• REST</li> <li>• OData</li> <li>• Open API (a.k.a Swagger)</li> <li>• GraphQL</li> </ul> </li> <li>• OAuth, OAuth 2.0 (JWT)</li> <li>• SAML</li> <li>• SSL</li> <li>• WS-Security</li> <li>• JMS listeners for message queues</li> <li>• Denodo Scheduler for batch process and lite ETL</li> </ul> <p><b>PERFORMANCE OPTIMIZATIONS</b></p> <ul style="list-style-type: none"> <li>• Dynamic Query Optimization for Analytics</li> <li>• Full and partial aggregation and join pushdown, even in federated views</li> <li>• Support for alternative data sources</li> <li>• On-the-fly data movement for optimization</li> <li>• Cost-based optimization (data statistics, data source indexes, data source execution model and parameters, network transfer rates)</li> <li>• Pushdown of selections/projections/ joins/groupby operations also on federated views</li> <li>• Multiple join strategies</li> <li>• Simplifying partitioned unions (Partition pruning)</li> <li>• Multi-mode caching: full, partial, incremental or total refresh, event based or scheduled, configured at the view level, incremental queries for SaaS sources</li> </ul>	<p><b>CACHE AND DATA MOVEMENT OPTIONS</b></p> <ul style="list-style-type: none"> <li>• Amazon Athena</li> <li>• Azure SQL</li> <li>• Azure SQL Data Warehouse</li> <li>• Azure Synapse Analytics</li> <li>• Amazon Redshift</li> <li>• Databricks 2.x</li> <li>• Delta Lake</li> <li>• IBM DB2 ( 8, 9, 10, 11 for LUW, 9, 10, 11 for z/OS)</li> <li>• Hive 2.0.0</li> <li>• Impala</li> <li>• MS SQL Server (2000, 2005, 2008, 2008R2, 2012, 2014, 2016, 2017)</li> <li>• MySQL (4 and 5)</li> <li>• Netezza (6 and 7)</li> <li>• Oracle (8i, 9i, 10g, 11g, 12c, 12c in-memory, 18c, 19c)</li> <li>• Oracle TimesTen 11g</li> <li>• PostgreSQL (9 and 10)</li> <li>• Presto</li> <li>• SAP HANA</li> <li>• Snowflake</li> <li>• Spark (1.5,1.6 and 2.x)</li> <li>• Teradata (12, 13, 14, 15, and 16)</li> <li>• Vertica (7 and 9)</li> <li>• Yellowbrick</li> <li>• Configurable “generic” adapter for other databases with JDBC drivers</li> </ul> <p><b>DATA PIPELINES</b></p> <ul style="list-style-type: none"> <li>• Remote Tables (created through UI or stored procedure)</li> <li>• Denodo Scheduler</li> </ul> <p><b>DATA GOVERNANCE</b></p> <ul style="list-style-type: none"> <li>• Data source refresh, change impact analysis, dependency tree, full data lineage</li> <li>• Denodo Governance Bridge: integration with IBM Information Governance Catalog</li> <li>• API to publish metadata including lineage information to data governance tools</li> </ul> <p><b>SECURITY</b></p> <p><b>Data in Motion – Secure Channels</b></p> <ul style="list-style-type: none"> <li>• Using SSL/TLS</li> <li>• Client-to-Denodo and Denodo-to-source</li> <li>• Available for all protocols (JDBC, ODBC, ADO.NET and WS)</li> </ul>	<p><b>Data at Rest - secure storage</b></p> <ul style="list-style-type: none"> <li>• Cache: third party database. Can leverage its own encryption mechanism</li> <li>• Swapping to disk: Serialized temporarily stored in a configurable folder that can be encrypted by the OS</li> </ul> <p><b>Encryption/Decryption</b></p> <ul style="list-style-type: none"> <li>• Support for custom decryption for files and web services</li> <li>• Transparent integration with RDBMs encryption</li> <li>• Encrypted metadata import/export</li> </ul> <p><b>User and Role Based including integration with AD/LDAP</b></p> <ul style="list-style-type: none"> <li>• Row and Column level authorization</li> <li>• Masking</li> <li>• Custom policies for specific security constraints and integration with external policy servers</li> </ul> <p><b>Authentication</b></p> <ul style="list-style-type: none"> <li>• Native and LDAP/Active Directory based Support for Kerberos and Windows SSO</li> <li>• Base64</li> <li>• Kerberos</li> <li>• NTLM</li> <li>• OAuth, OAuth 2.0 (JWT)</li> <li>• SAML</li> <li>• Two-factor authentication (through supported identity providers: Okta, Duo, etc.)</li> <li>• SSL</li> <li>• WS-Security</li> <li>• Pass-through session credentials</li> <li>• Leverage existing source privileges</li> </ul> <p><b>DATA MODELING</b></p> <ul style="list-style-type: none"> <li>• Design Studio: Web-based development studio for data modeling (desktop version still available)</li> <li>• Bottom-up and top-down (through interface views)</li> <li>• Denodo Model Bridge: Integration with third-party modeling tool <ul style="list-style-type: none"> <li>• ER/Studio Data Architect</li> <li>• ERwin Data Modeler</li> <li>• IBM InfoSphere Data Architect</li> <li>• SAP PowerDesigner</li> </ul> </li> </ul>
--	---	---

### 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 Query Monitor, Denodo Monitor Reports: all monitoring information is delivered via Denodo tools for real-time and historical analysis. Historical dashboards can be created using the monitor information
- Monitoring is also available via SNMP and JMX standards. Therefore interoperable 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
- Defining 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.
  - Installation of Denodo updates on clusters with several servers, without downtime.
  - Enable SSL in the Denodo servers.
- Subversion
- Microsoft TFS
- Git
- Resource Manager to limit and allocate resources to each session, role or user in a way that optimizes resource utilization for each application

- Change resource priority
- Enforce limited timeouts or limits on number of rows
- Add daily quotas per minute/day/month: e.g. only 50 queries per day

### USER INTERFACES

- Central Web Console: integrated access to all Denodo UIs with SSO (Kerberos, SAML, OpenID and OAuth)
- Solution Manager Web-based UI: 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)

