




The Denodo Platform


The AI Data Layer for Trusted Context and Enterprise Data Sharing

Enterprise data is increasingly distributed across cloud platforms, lakehouses, operational systems, SaaS applications, APIs, and legacy infrastructure. At the same time, business users, applications, and AI agents need faster access to **trusted data** and the **business context** required to understand and use it correctly.

Traditional architectures that depend primarily on moving and consolidating data cannot support every enterprise requirement. Replication can introduce latency, duplicate governance efforts, increase cost, and separate data from the operational systems where current business conditions are created. For analytics and AI, access to more data is not enough. Organizations need **live, governed, high-performing, and semantically consistent context** that can be applied wherever decisions and actions occur.

The Denodo Platform is the **AI Data Layer** that creates and operationalizes **active context** across distributed enterprise environments. Through a logical approach to data management, Denodo connects data across sources without requiring unnecessary movement, unifies business meaning through a shared semantic layer, and enforces governance and security as data is accessed. To support demanding enterprise workloads, Denodo applies a range of optimization and acceleration techniques — such as query pushdown, parallel execution, intelligent caching, Smart Query Acceleration, and lakehouse acceleration — so organizations can deliver trusted context with the performance their users, applications, and AI systems require.

Built on **universal connectivity, zero-copy delivery, governed access**, and **semantic trust**, the Denodo Platform turns distributed data into governed, reusable enterprise data products. Organizations can deliver **trusted data** faster, reduce integration and replication costs, apply governance consistently, optimize performance across distributed environments, and improve reuse across analytics, applications, enterprise sharing, and AI. This helps them modernize their data architecture while continuing to leverage existing investments.



Architecture for Modern Data and AI

To support this approach, the Denodo Platform focuses on the architectural capabilities required to deliver trusted data and business context to the people, applications, and AI systems that use it.

For many organizations, the biggest challenge is no longer storing or processing data at scale. It is delivering the right data, with the right business context, to the right consumer in a consistent, governed, and performant way. Data teams often spend significant effort building pipelines between systems, replicating datasets for different tools, tuning redundant copies for performance, and maintaining business logic across multiple environments. This increases operational complexity and makes it harder for analysts, developers, applications, and AI systems to access reliable data they can trust.

The Denodo Platform addresses this challenge by providing a unified data access, context, and delivery layer that simplifies how data is integrated, governed, shared, and consumed across the enterprise.

This architecture is built on five core capabilities:

LIVE DATA DELIVERY

Many modern use cases require access to current operational data and business context rather than periodically replicated datasets. Denodo enables real-time access to distributed data sources so analytics platforms, applications, and AI systems can work with up-to-date information.

Instead of relying solely on large-scale data movement, Denodo can query and combine data across systems at runtime while optimizing execution through techniques such as pushdown processing, parallel execution, intelligent caching, and selective materialization. This helps organizations reduce unnecessary replication, simplify pipeline management, and deliver fresher data for analytics and AI workloads without assuming that every request must be executed in the same way.

SEMANTIC CONSISTENCY AND SHARED BUSINESS CONTEXT

As data environments grow, business concepts are often defined differently across systems. Analysts, developers, applications, and AI models may interpret data inconsistently depending on the source.

Denodo provides a universal semantic layer that standardizes business definitions, metadata, relationships, metrics, and context across distributed sources. This enables teams to centralize transformation logic, maintain consistent definitions across tools, and improve data discovery and understanding. Capabilities such as metric views help organizations define, manage, and reuse trusted business metrics as part of this shared context, so analytics users, applications, and AI systems can work from consistent measures of the business.

Denodo also helps enrich the context around data products and other marketplace assets. Marketplace asset extensions enable organizations to represent additional relationships and business-relevant connections across assets, including

ontology-driven relationships where appropriate, helping users and AI systems better understand how data products, metrics, reports, policies, and related resources fit together. The result is greater semantic trust, with business users, applications, and AI systems working from a shared understanding of enterprise data.

CENTRALIZED GOVERNANCE AND SECURITY

Broader data access requires strong governance. Denodo enforces security and governance policies centrally while enabling data to remain in existing systems.

The platform supports fine-grained access controls, lineage tracking, and integration with enterprise identity and governance tools. Because these controls are enforced in the runtime data-access path, organizations can apply consistent security policies as data is requested and consumed, track data usage across analytics, sharing, and AI systems, and enable self-service access without compromising governance.

COST AND PERFORMANCE OPTIMIZATION

A logical, zero-copy approach does not mean compromising performance. Denodo includes multiple optimization and acceleration capabilities designed to deliver responsive access across distributed environments while helping organizations manage infrastructure and compute costs.

Denodo optimizes query execution through pushdown processing, parallel execution, cost-based optimization, intelligent caching, and selective data movement or materialization when it improves performance. Smart Query Acceleration can automatically recommend summaries to accelerate analytical workloads, while caching strategies help organizations balance freshness, performance, and cost based on workload requirements.

For lakehouse environments, Denodo also provides the Lakehouse Accelerator, powered by the Velox execution engine, which improves performance when accessing open

data lake formats such as Parquet, Delta, and Iceberg.

Together, these capabilities help organizations support high-performance analytics, applications, and AI workloads while reducing unnecessary replication, avoiding redundant pipelines, and optimizing execution across the enterprise data landscape.

BROAD ACCESS ACROSS THE ENTERPRISE DATA LANDSCAPE

Critical data rarely resides in a single platform. Operational systems, SaaS applications, partner environments, APIs, data lakes, and lakehouses all contain valuable information.

Denodo connects across this distributed landscape and exposes governed data and context through standard interfaces such as SQL, APIs, event streams, and Model Context Protocol (MCP). MCP support is especially important for agentic AI use cases, because it helps AI agents access governed enterprise context through a standard interface while preserving the security, semantics, and policy enforcement provided by the AI data layer.

This enables BI tools, applications, data scientists, developers, and AI agents to access governed data and context using the interfaces they already rely on.

Enabling the Last Mile of Data Delivery

These five capabilities address the final step in modern data architectures: delivering trusted, usable data and context to the people and systems that depend on it.

Denodo provides a consistent access layer that simplifies how data is discovered, integrated, governed, optimized, shared, and consumed across the enterprise. This enables organizations to deliver curated datasets for analytics, reusable data products for business domains, governed metrics for consistent decision-making, and reliable, context-rich data access for applications and AI systems.

These governed, semantically trusted data products provide a reusable and consumable expression of enterprise context for authorized users, applications, and AI agents. By enriching these assets with business definitions, relationships, metrics, lineage, governance policies, usage context, and performance optimization capabilities, Denodo helps organizations make active context available at the point of consumption without forcing every consumer or AI agent to recreate business meaning or integration logic independently.

By simplifying the last mile of data delivery, Denodo helps organizations modernize their data infrastructures while continuing to leverage existing investments in lakehouses, warehouses, operational systems, SaaS applications, and AI platforms.

The result is an infrastructure that delivers trusted data and business context consistently, securely, and efficiently to business users, data teams, applications, and AI systems.

Flexible Deployment and Scalable Editions

The Denodo Platform supports multiple deployment models to align with enterprise infrastructure strategies. Organizations can deploy Denodo in their own environments — on-premises, in private cloud, or within their public cloud accounts — or choose Denodo's fully managed cloud service.

Denodo is also available in multiple platform editions designed to support different workloads and adoption stages. Teams can begin with a focused deployment and expand over time as additional data domains, users, data products, metrics, AI workloads, and enterprise sharing requirements are added.

The table below outlines the capabilities available across the different Denodo Platform editions.

	Team Empower small teams to deliver focused data projects with fast time-to-value	High Availability Support growing workloads with reliable performance across multiple projects and teams	Business Critical Unify data at full enterprise scale with resilient, mission-critical operations
Configurations			
Deployment	Single Server ^[1]	Clustering	Clustering
Non-production environments	1 Development and 1 Staging	Unlimited	Unlimited
Disaster recovery environment	1	1	2
Solution Manager instances (Included)	1	1	1
Maximum cores	8	16	48
Data products included	100	225	750
Data volume included ^[2]	7.5 TB/year	25 TB/year	75 TB/year
Data sources	Unlimited	Unlimited	Unlimited
Data consumers	Unlimited	Unlimited	Unlimited
Features			
Data Marketplace users included	✓	✓	✓
Data Marketplace	✓	✓	✓
Dataset collaboration (endorsements, warnings, and deprecation notes)	✓	✓	✓
VQL procedures	✓	✓	✓
Denodo AI Assistant	✓	✓	✓
AI SDK	✓	✓	✓
Global security policies	✓	✓	✓
Smart Query Acceleration	✓	✓	✓
FinOps logging/integrated dashboard	✓	✓	✓
Advanced diagnostic and monitoring tool	✓	✓	✓
Data Lakehouse Accelerator (cores in MPP cluster) ^[3]	✓ (up to 512 cores)	✓ (up to 1024 cores)	✓ (up to 3072 cores)
Integration with VCS	-	✓	✓
Integration with 3rd party data modeling tools	-	✓	✓
Import governance tags from external catalogs	-	✓	✓
Clustering	-	✓	✓
Support			
Support level (included)	Standard	Premium	Global
Success services (included)	Enterprise Basic	Enterprise Advanced	Enterprise Advanced Plus
Onboarding services (not Included)	Basic	Advanced	Advanced

1. The Team tier is single server only, but it can be deployed in an active-passive failover configuration.

2. All data management units (KB, MB, GB, TB) used in this table and for billing purposes follow the International System of Units (SI) decimal standard where 1 unit = 1,000 of the preceding unit. Binary units (e.g., 1,024) do not apply.

3. Denodo offers an optional support package for customers needing technical assistance with the Data Lakehouse Accelerator.

Capabilities

DATA SOURCES

Relational Databases

- Generic (JDBC)
- IBM DB2 8.2, 9, 10, 11, and higher; 9, 10, and 11 for z/OS
- Denodo Virtual DataPort - 8.0, 9
- Apache Derby 10 and higher
- Informix 7, 12
- Microsoft SQL Server 2016, 2017, 2019, 2025, and higher
- MySQL: 4, 5, 8, and higher
- Oracle: 8i, 9i, 10g, 11g, 12c, 12c In-Memory, E-Business Suite 12, 18c, 19c, 23ai, and higher
- PostgreSQL 11, 12, 13, 14, 15, 16, 17
- Sybase ASE / SAP ASE 12, 15
- MS Access
- Microsoft Dataverse
- Huawei GaussDB

In-Memory Databases

- SAP HANA 1.0, 2.0
- Oracle TimesTen 11g
- Oracle 12c In-Memory

Parallel Databases and Appliances

- Apache Doris
- ClickHouse
- Exasol 7.1
- Greenplum 4.2
- Vertica 7, 9
- Netezza 4.6, 5.0, 6.0, 7.0, and higher
- Oracle Exadata
- ParAccel 8.0.2 (by using ParAccel 2.5.0.0 JDBC3g with SSL driver)
- SybaseIQ
- SQreamDB
- Teradata 12, 13, 14, 15, 16, 17, 20
- Yellowbrick

Cloud Data Warehouse / RDBMS

- Alibaba ApsaraDB for OceanBase MySQL
- Alibaba ApsaraDB for OceanBase Oracle
- Alibaba ApsaraDB RDS for MySQL
- Alibaba ApsaraDB RDS for PostgreSQL
- Alibaba ApsaraDB RDS for Microsoft SQL Server
- Alibaba ApsaraDB PolarDB for MySQL
- Alibaba ApsaraDB PolarDB for PostgreSQL
- Alibaba ApsaraDB AnalyticDB for MySQL
- Alibaba ApsaraDB AnalyticDB for PostgreSQL
- Alibaba MaxCompute
- Amazon Redshift
- Amazon Athena
- Amazon Aurora (MySQL and PostgreSQL)
- Amazon DynamoDB
- Azure Cosmos DB
- Azure SQL Database
- Azure Synapse SQL (aka Azure SQL Data Warehouse)
- Delta Lake

- Google AlloyDB
- Google BigQuery
- Google Cloud SQL for MySQL; Google Cloud SQL for PostgreSQL
- Google Spanner (Includes support for the service Spanner Data Boost)
- MongoDB Atlas
- Snowflake

Big Data

- Hive 0.13.0 (Hive Server 2); 1.1.0 (Hive Server 2); 2.0.0 (Hive Server 2); 3.1.2 and higher (Hive Server 3); Hive for Cloudera 1.1.0; Hive for Hortonworks 1.2.1
- Impala 1.2.4., 2.3, 3.x kudu
- Spark SQL 1.5, 1.6, 2.x, 3.x and higher
- PrestoDB
- Trino
- Databricks

NoSQL

- Amazon DocumentDB
- MongoDB
- Cassandra 3.x

Multi-Dimensional Sources

- Multidimensional database (generic)
- Azure Analysis Services
- SAP BI 3.x and BW 7.x
- Mondrian: 3.x
- IBM Cognos TM1
- MS SQL Server Analysis Services
- Essbase

Data Lake Storage/Formats

- Amazon S3
- Azure Data Lake Storage Azure Blob Storage
- Google Cloud Storage
- Parquet/Avro
- Delta/Iceberg
- Huawei Object Storage Service (OBS)

Web Services

- SOAP
- REST (XML, RSS, ATOM, JSON)
- OData v2.0 and v4.0

Flat and Binary Files

- CSV, pipe-delimited, regular expression-parsed
- Microsoft Excel xls 97-2003
- Microsoft Excel xlsx 2007 or later
- Microsoft 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

- Query vector-enabled databases, specialized vector search systems, and relational sources together while applying consistent security, lineage, and semantics across platforms such as Snowflake, Oracle 23ai, PostgreSQL, SQL Server 2025, BigQuery, Databricks, and Azure AI Search.
- CMS, file systems, text
- Elasticsearch

Cloud, SaaS, Web Sources

- Adobe Analytics
- AWS
- Google Analytics
- Google Sheets
- Facebook
- LinkedIn
- Microsoft Azure Data Lake
- Microsoft Sharepoint
- Microsoft Dynamics 365 Business Central / Customer Engagement
- Marketo
- ServiceNow
- Salesforce
- Twitter
- Workday
- many more through configurable JSON and XML adapters

Active Directory as Source or Leveraging Security

- LDAP v3
- Microsoft Active Directory 2003, 2008

Streaming/Messaging systems

- Kafka
- MQSeries
- SonicMQ
- ActiveMQ
- Tibco EMS
- Other JMS compatible services

Semantic Repositories

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

Packaged Applications

- Native SAP integration through ABAP Core Data Services (CDS) views
- SAP ERP/ECC (BAPIS and tables)
- Oracle E-Business Suite 12
- Siebel
- SAS 9.4

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
- Model Context Protocol (MCP) support for governed access by AI agents
- OAuth, OAuth 2.0 (JWT)
- SAML
- SSL
- WS-Security
- JMS listeners for message queues
- Denodo Scheduler for batch process and lite ETL

DATA MARKETPLACE

Cataloging

- Customizable for seamless data discovery and exploration for business users
- Business context, metadata, and documentation
- Domain centric business properties management
- Business categories and tags
- Intelligent search with smart ranking of results

Governance

- Graphical lineage spanning Denodo and Consumer tools such as Tableau and PowerBI
- 360-degree graph of relationships across data products, spanning Denodo and external assets
- Integrated request management (access, changes, data quality issues, etc.)
- Endorsement of datasets, comments, warnings, etc.
- Usage statistics: who uses what data, when and how
- Data profiling information

Self-Service

- A conversational interface for data discovery, data access, and deep research
- Context- and content-based search for finding relevant data products, metrics, reports, and related assets
- Graphical query for metric views, including visualization with Denodo Assistant
- Conversational interface in the SQL shell for building, explaining, and executing complex queries
- Last-mile data preparation wizards for customizing datasets by non-technical users
- Share queries or publish them as new data products
- Save personal queries for easy access
- Export to CSV, Excel, and Tableau Data Extracts
- Automatically match the UI language to browser settings

PERFORMANCE OPTIMIZATIONS

- Smart Query Acceleration for analytics
 - Aggregate Aware Summaries
- Lakehouse Accelerator integration for query acceleration and caching
- Support for Arrow Flight SQL, enabling high-performance data transfer
- Full and partial aggregation and join pushdown, even in federated views
- Support for alternative data sources

- On-the-fly data movement for optimization
 - Option to restricted data movement for sensitive data
- 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

DATA MATERIALIZATION OPTIONS

- Multi-mode selective data materialization (i.e. cache): full, partial, incremental, or total refresh, event-based or scheduled, configured at the view level, incremental queries for SaaS sources
- Directed Acyclic Graph (DAG) cache management automatically analyzes table dependencies and sequences replication tasks for optimal performance
- AlloyDB for PostgreSQL, Amazon Athena
- Supports schema evolution without data loss
- Amazon Aurora MySQL
- Amazon Aurora PostgreSQL
- Amazon Redshift
- Microsoft, Azure SQL
- Microsoft, Azure Synapse SQL (previously known as Azure SQL Data Warehouse)
- Clickhouse
- Databricks (with Apache Arrow support for accelerated pushdowns)
- Exasol
- Google BigQuery
- Google Cloud SQL for MySQL
- Google Cloud SQL for PostgreSQL
- Hive 2.0.0; Hive 3.1.2, and higher (HiveServer2)
- IBM DB2 (8, 9, 10, 11, and higher for LUW; 9, 10, 11 for z/ OS)
- Impala 2.3; 3.x Kudu
- Microsoft Fabric Data Warehouse
- Microsoft SQL Server (2000, 2005, 2008, 2008R2, 2012, 2014, 2016, 2017, 2019, and higher)
- MySQL 4, 5, 8, and higher
- Netezza 6, 7, and higher
- Oracle 8i, 9i, 10g, 11g, 12c, 12c In-Memory, 18c, 19c, and higher
- Oracle TimesTen 11g
- PostgreSQL 9, 10, 11, 12, and higher
- Presto
- SAP HANA 1 and 2
- Snowflake
- Spark SQL 2.x, 3.x, and higher
- Teradata 12, 13, 14, 15, 16, and 17
- Sybase
- Trino 4xx
- 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
- VQL stored procedures
- Support ETL and ELT patterns, depending on source and target capabilities, for use cases that require physical data movement or transformation.

LAKEHOUSE ACCELERATOR

- Presto-based MPP query engine leveraging the Velox C++ execution framework for high-performance access to data lake storage
- Graphical introspection of object storage (S3, ABFS, 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 enterprise governance catalogs, including Collibra Data Catalog, IBM InfoSphere Information Governance Catalog, and Microsoft Purview. OpenLineage support for metadata and lineage sharing
- 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
- Dynamic Access Control Policies
- Column and row restrictions, multiple masking options, deny execution

Authentication

- Native and LDAP/Active-Directory-based Support for Kerberos and Windows SSO
- 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

ADVANCED SEMANTICS

- Global security policies
- View- and column-level classification tags
- Support for the importing of external data governance tags from Collibra and other external data catalogs

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 tools:
 - ER/Studio Data Architect
 - ERwin Data Modeler
 - IBM InfoSphere Data Architect
 - SAP PowerDesigner
 - Sparx Systems Enterprise Architect

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 data quality 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.
- 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)

DENODO ASSISTANT

Data Marketplace

- A chatbot enabling users to discover data, explore views, and understand complex relationships
- Natural language queries with GenAI
- Support for OpenAI, Azure OpenAI, AWS Bedrock, and customer LLMs
- Personalized dataset recommendations
- Smart SQL-fragment autocomplete based on previous activity
- Multi-step conversation and disambiguation
- User profile with custom instructions
- Support for deep research and report generation
- Trusted data with exposed execution plan and methodology

Design Studio

- Summary recommendations for Smart Query Acceleration
- View and column description suggestions
- Business-friendly column name suggestions
- Data classification and tag suggestions
- AI-generated recommendations for join conditions
- Query diagnosis and query explanations
- Natural-language metadata discovery
- Natural-language data queries with SQL generation and execution options
- Ambiguity resolution and multi-step support for complex conversational workflows

VQL Function

- Automatically summarize, extract, analyze sentiment, remove sensitive data, and translate text using LLMs.

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 Denodo servers.
 - Update the Denodo version
 - Enable SSL in the Denodo servers.
- Multi-user development with version control integration
 - Ingratiation with source management systems such as Git
 - Workspaces for isolated development and testing

- 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:
 - 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 multilayered
 - Containerization support through Docker
- Public cloud:
 - Denodo Platform for AWS
 - Denodo Platform for Azure
 - Denodo Platform for Google Cloud
 - Denodo Platform for Alibaba Cloud
- Auto-scaling support both in AWS and Azure
- Agora - the Denodo Cloud Service

OPERATING SYSTEMS

- **Microsoft:** Windows Server 2022, Windows Server 2019, Windows Server 2016, Windows 11, and Windows 10
- **Linux:** Amazon Linux 2023, Ubuntu 20.04 LTS or later
- CentOS 9.x and 8.x, Red Hat Enterprise Linux (RHEL) 9.x and 8.x, Oracle Linux 9.x and 8.x, SUSE Linux Enterprise 15.x

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.



Denodo is the AI data layer that powers trustworthy agents and applications. The award-winning Denodo Platform enables that layer, transforming enterprise data into reliable insights for analytics and self-service. Organizations worldwide use Denodo alongside their data lakehouses to deliver AI-ready, business-ready data in a fraction of the time, achieving up to 4x faster time-to-insight, 345% ROI, and 10x better performance.

Visit www.denodo.com | Email info@denodo.com | Discover community.denodo.com

