



## The Denodo Platform for Google Cloud (GCP)

### Cloud data virtualization for faster, more agile data management

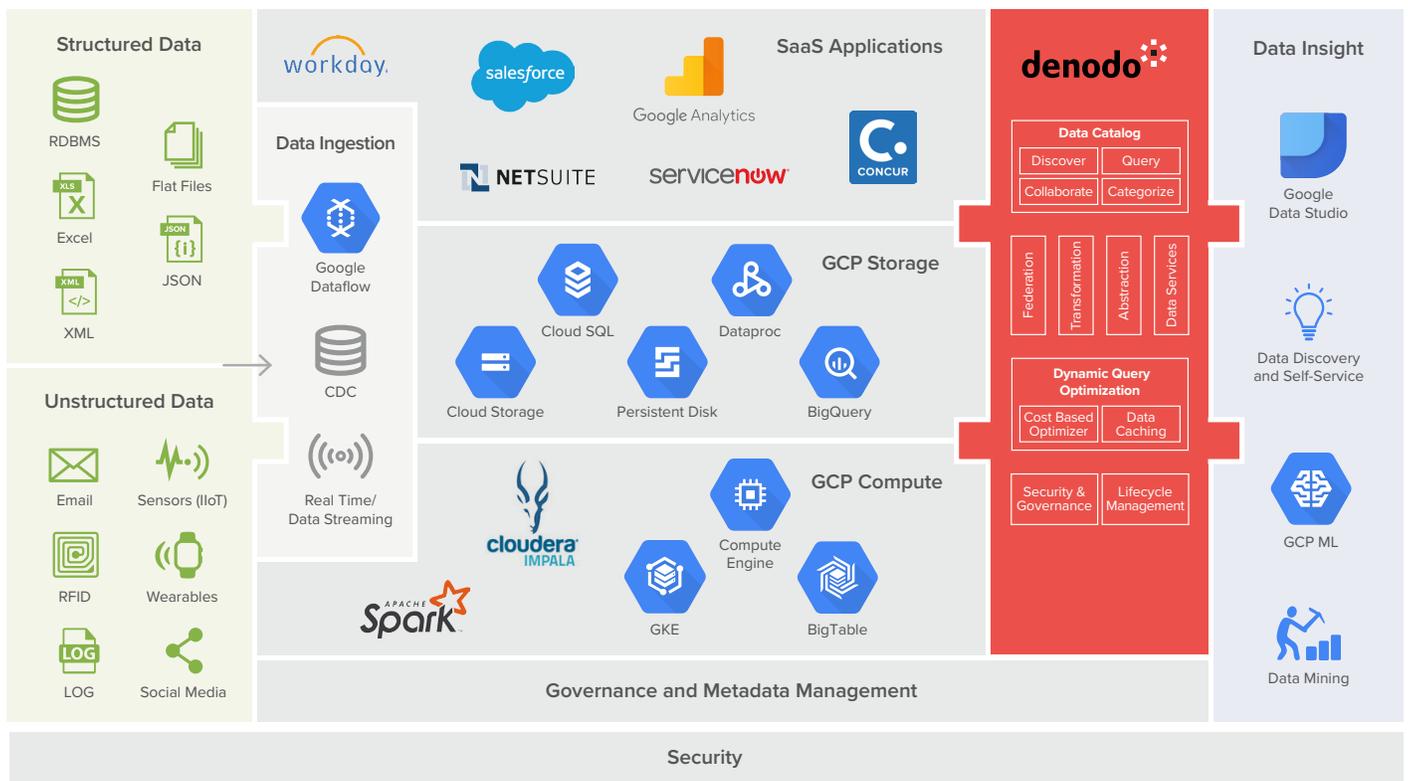
The Denodo Platform for the Google Cloud Platform (GCP) provides an enterprise-grade cloud data management solution that enables real-time data integration between on-premises, cloud, and multi-cloud data sources with all GCP data sources – Google BigQuery, CloudSQL, Cloud BigTable, Cloud Storage, Cloud DataProc, and more, including SaaS applications. The Denodo Platform offers agile, real-time access to a wide variety of data sources - structured and unstructured - providing business users with a faster, more semantically consistent foundation for cloud analytics and machine learning/AI.

The Denodo Platform provides easy-to-use tools that enable data engineers and integrators to quickly prepare integrated data sets for analysts to use in their GCP-based analytical sandboxes, while its integrated data catalog enables data stewards to explore, tag, and categorize data assets for easy discovery.

Leveraging Google’s flexible, rent-by-the-hour licensing, the Denodo Platform for GCP is offered at a wide range of pricing options including paying by the number of data sources. The Denodo Platform for GCP offers the same features as the award-winning on-premises version of the Denodo Platform.

The Denodo Platform for GCP accelerates performance in big data scenarios by leveraging massively parallel processing (MPP) processing capabilities with support for Spark, Impala, and Presto. The Denodo Platform for GCP simplifies the implementation of hybrid and multi-cloud data architectures and can be used as a virtual cloud data gateway to enable advanced analytics use cases such as analytical sandboxes, customer 360° analytics, and cloud data catalogs.

### Denodo Platform for GCP Reference Architecture



## Licensing and Deployment Options on GCP

The Denodo Platform for GCP can be deployed in the Google Cloud to support a distributed information architecture. It can be deployed via one of three options, and customers can choose the one that best meets their needs, depending on a variety of factors such as data-source location, the number of users and consuming applications, and the volume and type of queries.

- Using cloud Infrastructure-as-a-Service (IAAS) and leveraging BYOL (Bring your Own License). This option is best for customers with complex architectural deployments. Licensing is based on the number of cores on which the Denodo Platform is installed.
- Using the GCP Cloud Marketplace. This option provides pay-as-you-go flexibility and self-serviceability. This option provides multiple offerings based on the number of data sources that customers need to connect to, as shown in the table below.
- Microservices/Container deployment. The Denodo Platform is also available in Docker container format and can be installed with GCP containers and Kubernetes services (GKE) as needed.

## Benefits of the Denodo Platform on GCP

- **Easy access to a wide breadth of data:** Connect, combine, and consume data from a wide variety of data sources including Google BigQuery, Google Storage, and other Google cloud data sources, to provide a more data-driven approach to advanced analytics, AI, and machine learning.
- **Optimize access to big data:** Enhanced Google BigQuery support provides native connectivity and pushdown optimization to enable customers to process massive workloads with better query efficiency.
- **On-premises, cloud, and multi-cloud interoperability:** The Denodo Platform's ability to deliver real-time access to Google BigQuery cloud-native APIs enables frictionless data movement between existing on-premises and cloud data sources, and Google Cloud Storage.
- **Fast, intelligent execution of data queries:** BigQuery can be leveraged as a high-performance caching database for the Denodo Platform in the cloud. This supports advanced optimization techniques like data movement (data shipping), which enable multi-pass executions based on intermediate temporary tables.
- **Flexible packaging and pricing:** Users can leverage utility- based pricing through a number of offerings on the GCP Marketplace. They can start small and scale as needed.

## Did you know, the Denodo Platform is also available on the GCP Marketplace?

The [GCP Marketplace](#) offers convenience and flexibility:

- The marketplace provides a fast pathway to trying and buying software without having to go through sales negotiations or discussions about license terms and conditions.
- It offers one-click deployment to get started, along with help from the cloud support team.
- The 14-days Denodo Platform free trial via the marketplace provides a simple way to validate a use case.
- GCP marketplace offers several pricing options that are not easily available outside.
- Marketplace offerings also enable ease of scaling, and support peak-burst workloads such as month-end processing, enabling companies to scale up and down to save costs.

## Offerings Available on the GCP Marketplace

### UNRESTRICTED

This category has one option, unrestricted, which allows users to connect to an unlimited number of data sources without restrictions on the number of concurrent queries or returned results. This unleashes the full enterprise capabilities of the Denodo Platform for GCP.

### RESTRICTED BY DATA SOURCES

This category has four options, enabling users to select between two and ten data sources. For each of these options, there is no limit on the number of queries or result sets.

- 2 data sources
- 5 data sources

## Multiple Use-Case Support with the Denodo Platform on GCP

**Data-as-a-Service (DaaS):** Information agility and provisioning is a key data distribution model in which a variety of data (structured and semi-structured) across Google cloud storage and web services can be made available to users as an API endpoint. Denodo API support can help expose this data in a centralized manner, streamlining the DaaS model.



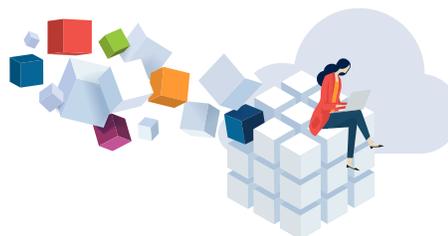
**Integrated Security for Cloud Sources:** Data virtualization acts as a gatekeeper to the cloud, extending cloud sources with limited security features. It enables single sign-on (SSO) for external sources and the integration of local users and groups (LDAP, Windows AD) with cloud sources that do not offer such capabilities by default. It also enables fine-grained security such as column- and row-level restrictions, and masking and integration with cryptographic appliances.



**Abstracting SaaS APIs in the Cloud:** Enables traditional reporting tools to work with any kind of SaaS API. For example, for SQL-to-SaaS processes, data virtualization abstracts SaaS APIs (usually REST services) as part of a relational model, and avoids replicating cloud data back to the data center.



**Machine Learning/Artificial Intelligence (ML/AI) and Data Science in the Cloud:** With the Denodo Platform for GCP, customers can leverage the Denodo Platform's data catalog to search and tag the right datasets for analytics and ML projects. This will also help data scientists to combine data stored in the virtual data lake and data virtualization layer to build models in a quick and easy manner, putting cloud elasticity to work, and using the data lineage capabilities of the Denodo Platform to access all of the data in a governed fashion.



**Simplifying Cloud Migration (Data warehouse modernization or data lake creation):** Data virtualization acts as a common access point for both internal and external data sources, data warehouse modernization using Google BigQuery, and Google DataProc, providing a single schema with no replication: Virtual data lakes enable the combination of data across sources, regardless of type or location, enabling the definition of a common semantic model across disparate sources.

## Key Data Sources Supported in the GCP Cloud

- Google BigQuery, CloudSQL, CloudSpanner – Easily connect using the BigQuery JDBC driver and generic JDBC drivers.
- Google Cloud DataProc – Cloud DataProc is Google’s Apache Hadoop and Apache Spark service running on GCP. It supports the standard Hadoop components and interfaces, such as HBase, Spark SQL, and HDFS.
- Google Cloud Storage and Bigtable – Access data stored in cloud storage directly from the Denodo Platform. Access CSV, JSON, XML, delimited, and Parquet files stored in storage buckets. The Denodo Platform has a special connector for directly accessing object storage files.
- Cloud Dataflow, Snowflake – Cloud Dataflow integrates with JMS and is supported with Denodo. Snowflake is a cloud-based data warehouse as a service (DWaaS) offering on GCP. Connect to Snowflake using the Snowflake JDBC driver.
- A Wide Variety of SaaS Data Sources – These include Salesforce, ServiceNow, and others.

## Cloud-Specific Features and Optimizations in the Denodo Platform for GCP

- Specific connectors for cloud-based databases (Google Big Query, BigTable, Cloud Storage, etc.) and native bulk-data load techniques for query optimization and caching support.
- Highly configurable vendor- and version-agnostic connectors for REST and OData, enabling connection to any HTTP API.
- Support for different formats (XML, JSON, CSV) and HTTP methods (GET, POST), Native support for OAuth, NTLM, SPNEGO, etc.
- Native SFDC adapter with full SOQL support, taking advantage of Salesforce-specific query language.
- Caching with incremental queries -- merge cached data and fresh data to provide fully up-to-date results with minimum latency.

## DENODO 14-DAY FREE TRIAL ON GOOGLE CLOUD

A **free 14-day trial** of the Denodo Platform for GCP, from the GCP Marketplace, is the fastest way to get started with a fully functional Denodo Platform environment. Start small and grow at your own pace. The Denodo Platform for GCP offers a variety of deployment options, such as unlimited data sources, depending on the scope and complexity of each customer’s environment. Technical support is included.

Those new to the Denodo Platform can also leverage the free Denodo **Test Drive** on GCP, which provides a quick and easy way to get started.

