



Expert Trail: Deployment Topology

Revision 20211109

NOTE

This document is confidential and proprietary of **Denodo Technologies**.
No part of this document may be reproduced in any form by any means without prior written authorization of **Denodo Technologies**.

Copyright © 2022
Denodo Technologies Proprietary and Confidential

CONTENTS

1 LOOKOUT.....	3
2 THE HIKE.....	4
3 EXPLORATION.....	6
4 GUIDED ROUTES.....	10
4.1 DENODO TRAINING COURSES.....	10
4.2 TECHNICAL ADVISORY SESSIONS.....	10
4.3 PROFESSIONAL SERVICES.....	11
5 BIG HIKE PREP CHECK.....	11

1 LOOKOUT

Expert trails guide Denodo users through all the relevant materials related to a specific topic, including official doc, KB articles, training, Professional Services offering, and more. The main goal is to give users a single place with references to all the information that they need to become a Denodo expert on any specific topic.

The Deployment Topology expert trail will provide you with the necessary knowledge for defining your Data Virtualization architecture, not only in the design phase, but also you will get expert knowledge on how to implement it such as installing the Denodo Platform or the Solution Manager, configuring a cluster or deploying your machines in the cloud.

2 THE HIKE

Stage 1: Understanding your scenario

Before defining your architecture, you may need to get some inspiration by taking a look at some [Data Virtualization Use Cases](#). In addition to this section, you have dedicated sections for some of the most common use cases for Data Virtualization that include Customer Success stories and detailed solutions for each particular use case:

- [BI and Analytics](#)
- [Cloud Solutions](#)
- [Data Governance](#)
- [Data Services](#)
- [Master Data Management](#)
- [Denodo in Data Science and Machine Learning Projects](#)

A complete analysis of your use case should be performed:

- Which are my sources? Are they located in the cloud or on-premises?
- What kind of client applications are in place? How are they going to connect to Denodo? Are they cloud-based or on-premises?
- How many users are expected?

Stage 2: Defining your infrastructure architecture

Now it is time to decide on the infrastructure architecture. Depending on the analysis done in the previous stage, you may be interested in a pure on-premises architecture, cloud, or hybrid architecture.

The [Denodo Cloud Architectures](#) presentation provides a detailed overview of reference architectures for cloud environments.

In the [Denodo for Cloud](#) section, you can review the available options in the Cloud for installing the Denodo Platform. Bring your Own License or Cloud Marketplace out-of-the-box images are available for fitting your needs.

The [Multi-Layer Architecture](#) presentation provides information about how to set up a multi-layer architecture that facilitates moving processing to data location.

Either if you decide to go cloud with your Own License or on-prem, you can check the [Hardware Requirements](#) section of the Denodo Platform Installation Guide to make your decisions on the architecture design. The same section for the Solution Manager is also available: [Solution Manager Hardware Requirements](#).

You may be considering a containerized deployment in Docker. You can review the [Denodo Platform and Docker datasheet](#) and the [Denodo on Docker FAQ](#) to get more information

that will guide your decision.

Security decisions should be made at this point. You can take a look at the [Expert Trail: Security](#) for getting a detailed overview of the security features of the Denodo Platform.

Stage 3: Making it real!

Once you have decided on the architecture, it is time to make things real!

The first step on your new installation is installing the Denodo Solution Manager that will manage licenses for all your environments and will provide a set of tools to manage and administer your environments. The section “Denodo 8.0 Installation Keys” of the [Ready to work with Denodo 8.0](#) KB article is a good point to start since it explains the key steps for this installation process.

Once the Solution Manager is up and running, it’s time to install the Denodo Platform servers. Check the [Denodo Platform Installation Guide](#) for more details on these steps. In addition, you need to configure the [Cache Module](#) using an external database. In addition, Denodo 8 allows storing the metadata in an external metadata database, but this option is only recommended in some specific scenarios that can be checked in the [Storing the Metadata on an External Database](#) section of the VDP Administration Guide.

More details about this can be found in the [Expert Trail: Operations](#) and [Expert Trail: Cache](#)

Are you thinking about a cluster? Review the [Denodo Platform Cluster Architecture](#) to review the important points to keep in mind for building a Denodo cluster.

In case you are installing on a Cloud machine, you may be interested in reviewing the [Installation steps on a cloud environment](#) article to get some important steps to be performed on a cloud environment.

Instead of performing the installation by yourself, you can decide to use the available option in the Cloud Marketplaces. In that case, you can check the following links depending on the vendor you have selected to deploy the desired cloud instances for your environment:

- **AWS:** [Denodo Platform for AWS](#)
 - [AWS Marketplace link](#)
- **Azure:** [Denodo Platform for Azure](#)
 - [Azure Marketplace link](#)
- **Google Cloud:** [Denodo Platform for Google Cloud Platform](#)
 - [Google Cloud Marketplace link](#)

You have the option to select Denodo Platform or the [Denodo Standard](#) version. In addition, you can select the options where the license is paid per use or you can bring your own license (BYOL options).

The Automated Mode for AWS included in the Solution Manager allows automating the deployment of Denodo on AWS. If your selected Cloud vendor is AWS you can take a look at this [Quick Start Guide with Automated Cloud Mode](#).

Have you decided to go on Docker and Kubernetes? Review the [Deploying Denodo in Kubernetes](#) and the [Deploying the Solution Manager in Kubernetes](#) for more information on this scenario.

Stage 4: Monitoring your deployment

Once the environment is up and running, it's a good practice to monitor it to check if everything is working as expected.

The Denodo Platform provides different endpoints to perform a Health Check of the different components. These are described in the [Health Monitoring](#) section of the Denodo Platform Installation Guide.

At some point, you may be interested in analyzing if the deployment size is enough for the current workload or the expected load in the future. You can monitor the status of the server in real-time using the [Diagnostic and Monitoring Tool](#) or the [JMX interface](#) or by using the [Denodo Monitor](#) that will generate logs that can be later analyzed to get information on some important metrics such as CPU usage, memory, threads, requests, etc.

You can combine these Denodo tools that collect information about the status of the servers, with an external load testing tool to perform your Load Tests to ensure your server is ready for the current or expected load. You can get information on how to perform these load tests with some of the well-known tools in the [Denodo Load Testing with LoadRunner](#) or [Denodo Load Testing with Apache JMeter](#) KB documents.

More information about Monitoring can be found in the [Expert Trail: Monitoring](#).

3 EXPLORATION

Fill up your backpack with additional gear:

Use Cases

Official Documentation	<ul style="list-style-type: none"> ● Virtual DataPort General Architecture
KB Articles	<ul style="list-style-type: none"> ● Data Virtualization and ETL ● Denodo Cloud Architectures ● Multi-Layer Architecture ● Denodo in Data Science and Machine Learning Projects
Webinars	<ul style="list-style-type: none"> ● Bringing a real time flavor to your enterprise analytics ● Centralize Security and Governance with Data Virtualization ● Fast Data Strategy Virtual Summit 2021
Additional Resources	<ul style="list-style-type: none"> ● Albertsons Leverages the Denodo - Case Study Denodo

- [Toyota-Astra Motor Leveraged the Denodo Platform to Simplify its Data Landscape and Achieve a Single Version of the Truth - Case Study | Denodo](#)
- [St. James's Place Wealth Management - Case Study | Denodo](#)
- [Enabling a Customer Data Platform Using Data Virtualization - Analyst Report | Denodo](#)
- [Modernizing Data Architectures for a Digital Age Using Data Virtualization - Analyst Report | Denodo](#)

Cluster architectures

Official Documentation

- [Solution Manager - High Availability](#)
- [Cluster Architectures / Server Backup — Virtual DataPort Administration Guide](#)

KB Articles

- [How to configure a VDP Backup server in the Solution Manager](#)
- [Configuring a Denodo Cluster with BigIP F5](#)
- [Configuring a Denodo Cluster with HAProxy Load Balancer](#)
- [Installing the Data Catalog on a Cluster: Sharing the Same Settings across all the Nodes](#)
- [Installing the Information Self-Service Tool on a Cluster: Sharing the Same Settings Across all the Nodes](#)

Cloud Deployments

Official Documentation

- [Recommendations to Run the Denodo Platform on Amazon AWS](#)
- [Recommendations to Run the Denodo Platform on Microsoft Azure](#)
- [Creating A Custom AMI for the Denodo Platform — Solution Manager Administration Guide](#)

KB Articles

- [Denodo Cloud Architectures](#)
- [Configuring Auto Scaling of Denodo in AWS](#)

Webinars

- [Data Virtualization to Survive a Multi and Hybrid Cloud World](#)
- [Simplifying Your Cloud Architecture with a Logical Data Fabric](#)
- [Data Integration in the Cloud](#)

Additional Resources

- [Denodo Standard for Cloud Marketplaces - Cloud Howto Videos | Denodo Community Site](#)
- [Automatic creation of environments in AWS - Videos | Denodo](#)
- [Leveraging data virtualization in the Azure Cloud - Case Study | Denodo](#)

Containers - (Docker - Kubernetes)

Official Documentation	<ul style="list-style-type: none"> ● Running the Denodo Platform on Docker ● Denodo Platform Container Quick Start Guide
KB Articles	<ul style="list-style-type: none"> ● Deploying Denodo in Google Kubernetes Engine (GKE) ● Deploying Denodo in Amazon Elastic Container Service (ECS) using Fargate ● Deploying Denodo in Amazon Elastic Kubernetes Service (EKS) ● Deploying Denodo in the Azure Kubernetes Service (AKS) ● Data Persistence in Containers

Denodo Components Installation

Official Documentation	<ul style="list-style-type: none"> ● Denodo Platform Installation Guide — Installation Guide ● Solution Manager Installation Guide — Solution Manager Installation Guide ● Denodo Platform 8.0 Supported Operating Systems ● Supported JDBC Data Sources — Virtual DataPort Administration Guide ● Preparing the Connection to Data Sources — Installation Guide ● Default Ports Used by the Solution Manager Modules — Solution Manager Installation Guide ● Default Ports Used by the Denodo Platform Modules ● Denodo Security Token — Solution Manager Administration Guide ● Setting-up Kerberos Authentication — Installation Guide
KB Articles	<ul style="list-style-type: none"> ● Solution Manager Installation Checklist ● Denodo Platform Installation Checklist

Monitoring

Official Documentation	<ul style="list-style-type: none"> ● Monitoring the Virtual DataPort Server — Virtual DataPort Administration Guide ● Monitoring — Solution Manager Administration Guide
KB Articles	<ul style="list-style-type: none"> ● Automating status check of Denodo services ● Denodo Monitor as a Windows service ● Denodo Monitor as a Linux Service ● Denodo Monitor timings ● Monitoring Denodo with Azure Monitor ● Monitoring Denodo with Splunk ● Monitoring Denodo with Amazon CloudWatch ● Monitoring Denodo servers with JMX

Additional Resources

- [Monitoring the embedded Tomcat web container](#)
- [How to store Denodo logs in Amazon S3](#)
- [Configuring and starting the Denodo Monitor - Videos | Denodo Community Site](#)
- [Storing Denodo Monitor logs in an external Database - Videos | Denodo Community Site](#)
- [Launching the Denodo Monitor with Solution Manager - Videos | Denodo Community Site](#)
- [Denodo Platform Monitoring - Tutorial | Denodo Community Site](#)

4 GUIDED ROUTES

4.1 DENODO TRAINING COURSES

Denodo training courses provide expert data virtualization training for data professionals, including administrators, architects, and developers.

If you are interested in Deployment Topology you should enroll in the following course/s:

- **Data Virtualization Architect:** This training covers architectural patterns for the use of data virtualization and best practices for performance optimization, data services governance, scalability, and operation.
- **Denodo Deployment Configuration:** This course talks about architecting the infrastructure depending on the Denodo License restrictions.
- **Denodo Platform Endpoints:** This course covers all the northbound endpoints supported by Denodo and how to use them. This course can help to design the architecture based on the supported endpoints.

4.2 TECHNICAL ADVISORY SESSIONS

Denodo Customers with active subscriptions have access to request [Meet a Technical Advisory sessions](#).

These are the sessions available related to deployment topology.

Hardware & Software Requirements	Denodo Platform & Solution Manager Architecture Overview 	Overview of the Denodo Platform components architecture. Describe all the different Denodo Platform components (VDP Server, Scheduler, Solution Manager, Data Catalog, etc.): - Mission. - How they interact with each other. - How to interact with client applications and sources. - How they fit in your organization.
	Hardware & Software Requirements for Denodo Platform Components	Describe the Hardware and Software Requirements for the Denodo Platform: both the main installation and the Solution Manager installation. Assisting in: - Selecting the appropriate hardware resources and software for your organization (physical, VM, container). - Selecting cache, metadata storage, and VCS system.
	Deployment Sizing	Recommend Denodo Platform deployment sizing. This will be based on the current + expected workload.

		This session is complementary to the Capacity assessment which will provide accurate information on the current workload. Guide on sizing your solution according to your plans.
Infrastructure Architecture	Cluster Deployment Topology 	<ul style="list-style-type: none"> - Explain the Denodo Cluster architecture and best practices. - Advice on defining the cluster architecture for your deployment. - Review your current architecture design. - Define steps to configure the Load Balancer. - Deployment in containers (e.g. Docker).
	Cloud & Hybrid Deployment Topology 	<ul style="list-style-type: none"> - Explain the current offering of the Denodo Platform in the Cloud (AWS, Azure, GCloud). - Review the specific considerations when deploying in the cloud. - Explore the different Hybrid (Cloud - On-Prem) Architectures. - Discuss pros and cons and guidelines to determine the best solution for your case. - Use of Solution Manager Cloud Automated Mode.
Architecture Definition	Reference Architecture	Assistance in defining or reviewing your Reference Architecture. Where to position the Denodo Platform in your ecosystem and the main role that the Denodo Platform can play in your organization's use cases.

4.3 PROFESSIONAL SERVICES

Denodo Professional Services can help you at the start or any part of your query performance trail. You can find information about the Denodo Professional Services offering in:

[Professional Services for Data Virtualization | Denodo](#)

In particular, you may be interested in the following module:

- **Virtualization Architecture & Design**

Additional other related modules could be:

- **Use Case Discovery**
- **Denodo Capacity Assessment Module**

If you are a Denodo customer, you can reach out to your Customer Success Manager for details about any Guided Route that you need.

5 BIG HIKE PREP CHECK

Let's see if you are ready to start your big trail. Take this 2-question questionnaire to check your readiness for an enjoyable hike.

Read the questions below, think about the solution and check if you got them right by looking at the solution. Have you become an expert?

1. In our company, we have a traditional DataWarehouse based on an on-prem Oracle server containing a significant amount of the data of the company and an AWS Athena database located in the AWS public cloud with a significant amount of customer's data. We would like to implement a Logical Data Warehouse using Data Virtualization, which is the recommended Denodo architecture for implementing Denodo a solution for this scenario?

[Click here to check if you got it right](#)

The recommended approach for this scenario is a **Denodo Hybrid architecture** where two different Denodo servers can work together in a two-zoned architecture (cloud-on prem), one installed on-prem, next to the on-prem data sources like our Oracle server and another installed in the same cloud as the other sources, in this scenario in the AWS cloud. Applications in every zone will have access to the data from the other zone using Denodo that will reduce the amount of data transfer through the network. This will minimize the network overhead since aggregated data from the external sources is transferred across the network in both directions, in response to the consumer's requests. This scenario is explained in the "Denodo Hybrid Deployment (On-premise & Cloud)" section of the [Denodo Cloud Architectures](#) presentation described in Stage 1.

2. Our company has a big Denodo development team working on a Denodo 7 development single server where every developer has its own user and connects to the server for creating the required views. The architect has decided to migrate to Denodo 8 and also, based on new requirements, the usage of Denodo is planned to be increased so the development team will have 3 or 4 times the previous size, hence the load of the development server will grow due to new users working concurrently on the server. Keeping in mind this new scenario, how would you design the new development environment using Denodo 8?

[Click here to check if you got it right](#)

The new scenario requires a higher capacity for the Denodo 8 development environment. The recommended architecture for this scenario is a **cluster** of Denodo 8 servers in order to avoid that a higher number of developers executing concurrent queries impacts the development team. An **external metadata database** is recommended to store the server metadata shared among the different nodes. This is based on the recommendations described in the [Storing the Metadata on an External Database](#) section of the VDP Administration Guide, keeping in mind in the development environment, changes in elements are not done in a controlled way and this is one of the scenarios where the usage of an external metadata database is recommended.