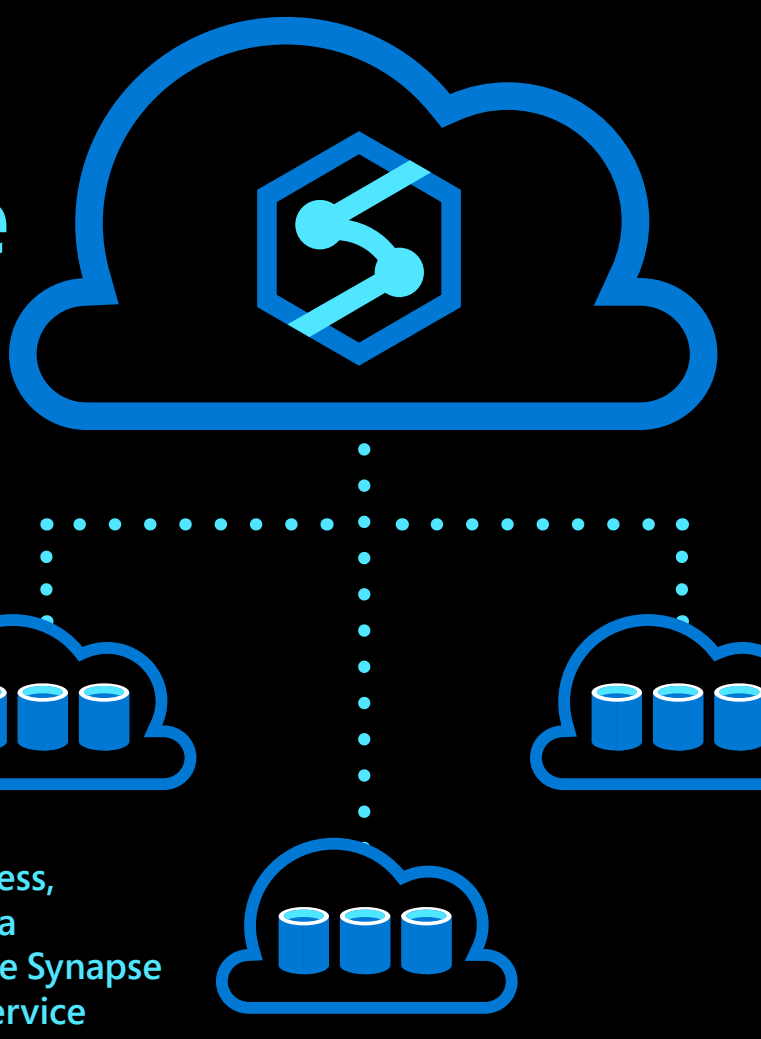


Azure Synapse Serverless SQL Pool 101



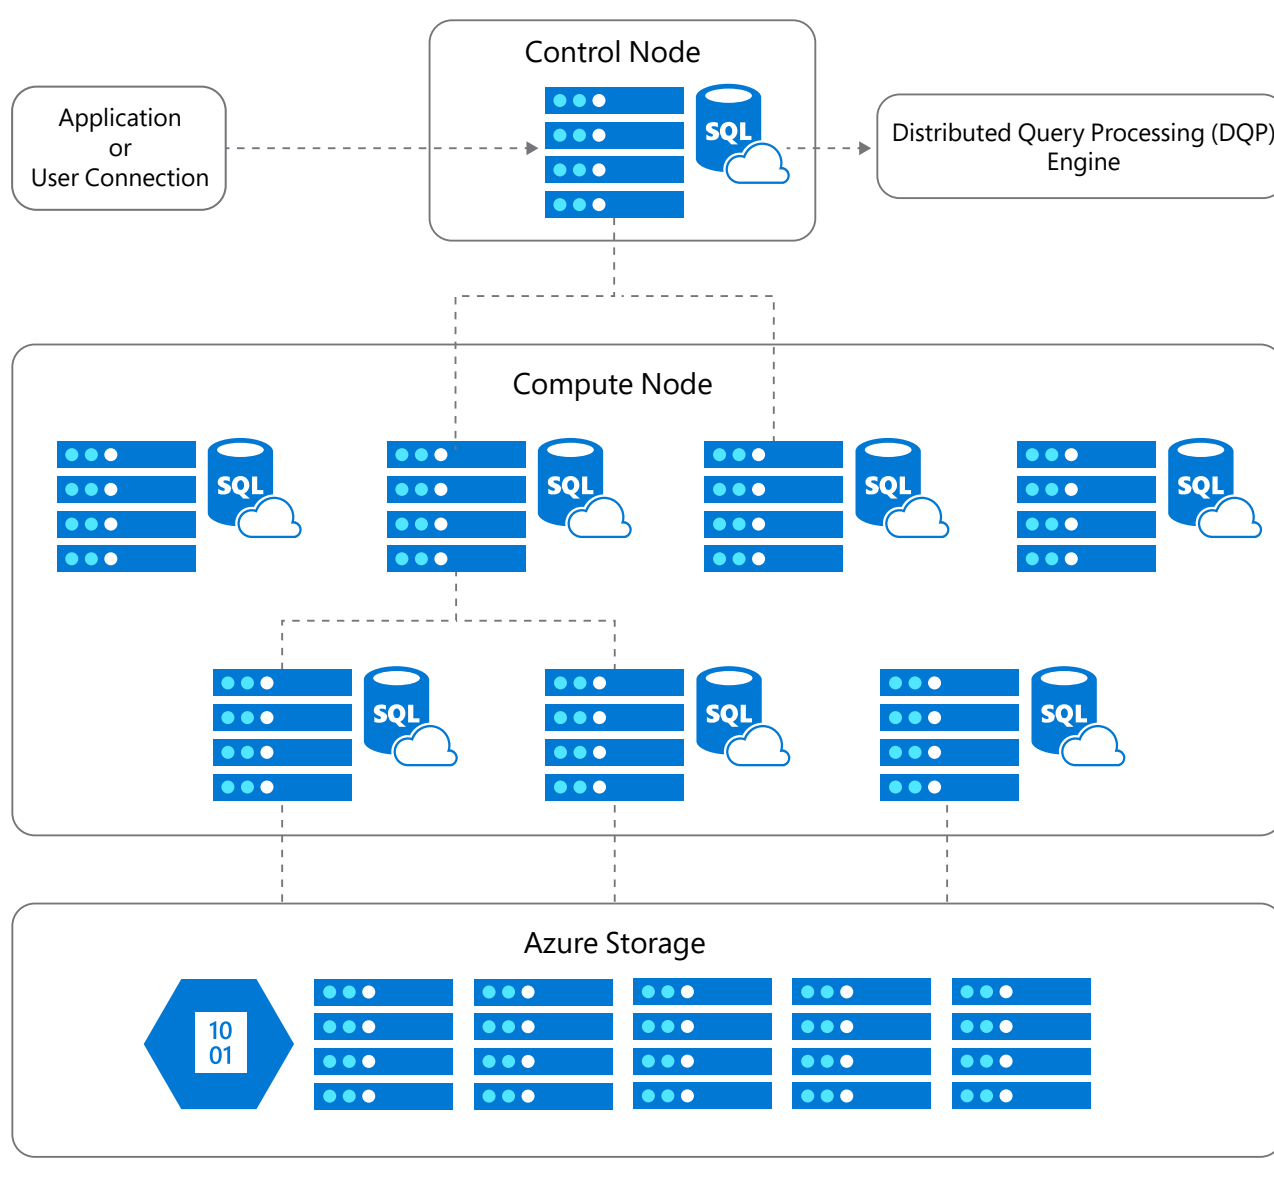
Azure Synapse serverless SQL is a serverless, distributed data processing service with a consumption-based pricing model. Azure Synapse serverless SQL pool is a pay-per-query service that enables you to run T-SQL queries on-demand without requiring you to provision any resources.

The anatomy of Azure Synapse serverless SQL pool

Azure Synapse serverless SQL pool employs a node-based structure and uses a scale-out architecture to distribute computational processing of data across multiple nodes. With decoupled storage and compute, you can benefit from being able to independently scale the compute power regardless of your storage requirements.

How does Azure Synapse Serverless SQL pool work?

1. An application connects and issues T-SQL commands to a **control node**, which is the single point of entry for Azure Synapse serverless SQL pool.
2. The **Distributed Query Processing (DQP) engine** runs on the **control node** to optimize and coordinate parallel queries.
3. This **control node** utilizes a **DQP engine** to optimize queries for parallel processing and then passes operations to **compute nodes** to do their computational work in parallel.
4. Azure Synapse serverless SQL pool allows you to query files in your **Azure Storage** accounts using the **DQP engine**. You can conveniently run T-SQL queries against your data on-demand without the need to ingest data into data stores or provision any resources.



Why should you use Azure Synapse serverless SQL pool?

Self-managed infrastructure
Just let the system automatically scale itself based on your query resource requirements.

Use it immediately
You can start querying your data lake files as soon as your Azure Synapse workspace is created.

True pay-per-use model
You only pay for the queries you execute.

Scenarios and benefits of using Azure Synapse serverless SQL pool

Scenario: Basic discovery and exploration
Quickly analyze the data in various formats (Parquet, CSV, JSON) in your data lake, so you can plan how to obtain insights from it.

Scenario: Logical data warehouse
Provide a relational abstraction on top of raw or disparate data, allowing an up-to-date overview without having to relocate or transform your data.

Scenario: Data transformation
Have a simple, scalable, and efficient way to transform data in the lake using T-SQL, so it can be fed to BI and other tools, or loaded into a relational data store (Azure Synapse SQL pool, Azure SQL Database, and so on).

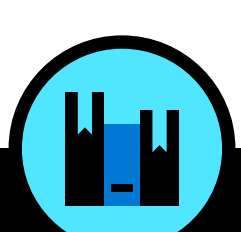
How does each professional role benefit from Azure Synapse serverless SQL pool?

Data Engineers
Can [explore the data lake, transform and prepare data, and simplify their data transformation pipelines.](#)

BI Professionals
Can quickly [create Power BI reports on top of data in the lake](#) and Spark tables.

Data Analysts
Can [explore data and Spark external tables](#) created by data scientists or data engineers using the familiar T-SQL language or their favorite tools, which can connect to Azure Synapse serverless SQL pool.

Data Scientists
Can quickly reason about the contents and structure of the data in the lake, thanks to features such as [OPENROWSET](#) and automatic schema inference.



Further learning and resources

- > [Azure Synapse Serverless SQL pool 101 e-book](#)
- > [Sign up for a free Azure account](#)
- > [Explore the Azure Synapse Analytics Knowledge center](#)
- > [Read the Cloud Analytics with Microsoft Azure e-book](#)
- > [Learn more about serverless SQL pool in Azure Synapse Analytics](#)
- > [Join the hands-on training series for Azure Synapse Analytics](#)
- > [Download the Azure Synapse Analytics proof of concept playbook](#)

Speak to a sales specialist for help with pricing, best practices, and implementing a proof of concept