

Get started with SQL Database in Microsoft Fabric

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/1-introduction>

Introduction

Completed

SQL Database in Microsoft Fabric is a powerful and versatile solution designed to integrate seamlessly with the broader Microsoft ecosystem. It allows users to use the capabilities of SQL databases within the Microsoft Fabric environment, providing a unified platform for managing and analyzing data. This integration enables organizations to streamline their data workflows, enhance collaboration, and drive more informed decision-making processes.

Moreover, SQL Database in Microsoft Fabric supports a wide range of use cases, from operational and transactional workloads to advanced analytics. This flexibility makes it an ideal choice for organizations looking to consolidate their data infrastructure and use the full potential of their data assets. With features like seamless integration with other Microsoft services, enhanced security, and scalability, SQL Database in Microsoft Fabric empowers users to build and deploy data-driven applications with ease.

The **Databases** workload in Fabric is designed to address many aspects of operational SQL databases, enabling data engineers, analysts, and data scientists to work together to create and query data that is optimized for their specific needs.

In this module, you'll learn about SQL database in Fabric, and how you can manage and query your data efficiently. You'll explain how SQL database in Fabric works, understand security options available, explain how to use Copilot to enhance troubleshooting, and explore mirroring and data virtualization capabilities.

2. Work with SQL databases

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/2-work-with-sql-database>

Work with SQL databases

Completed

SQL Database in Microsoft Fabric is a versatile and developer-friendly transactional database built on the foundation of Azure SQL Database. It allows for the creation and management of operational databases within the Fabric environment.


Differently than Azure SQL Database, which is a Platform as a Service (PaaS), SQL Database in Microsoft Fabric is a Software as a Service (SaaS). This means that users can enjoy a low-maintenance solution, allowing them to focus even more on their core business activities.

One of its capabilities is the automatic replication of data into OneLake and conversion to Parquet in near real-time, which facilitates analytics without the need for complex ETL processes. This integration ensures that data is always up-to-date and accessible for various services within Fabric, such as Spark for analytics, notebooks for data engineering, and Power BI for visualization.

Create a SQL Database

To create a new SQL database in Fabric, you need a new or existing workspace. Start by navigating to the Fabric portal and selecting **Databases**. Under the **New** section, select the SQL database tile. Enter a name for your new database and select **Create**.


Once the database is provisioned, you see the **Explorer** pane on the **Home** page displaying the database objects.



To help you get started, there are three useful tiles under **Build your database**. The **Sample data** option allows you to import the *AdventureWorksLT* sample data into your empty database. The **T-SQL** option provides a web-editor for writing T-SQL to create database objects like schemas, tables, and views. The **Connection strings** option displays the SQL database connection string needed for connecting with SQL Server Management Studio or other external tools.

Query a SQL Database

You can query a SQL database in Fabric using similar tools available for Azure SQL Database, with the added convenience of a web-based editor in the Fabric portal. This provides an end-to-end, integrated product that simplifies analytics and fosters collaboration.



The **Open in** option allows you to launch Visual Studio Code and SQL Server Management Studio (SSMS) with the connection properties prefilled, making it easier to connect and start working immediately.

Source control

Source control is an essential aspect of managing SQL databases in Microsoft Fabric. It allows you to track changes, collaborate with team members, and maintain a history of modifications. When you integrate your SQL database with a source control system, you ensure that all changes are documented and can be reverted if needed. This practice enhances the reliability and consistency of your database development process.

If you're familiar with source control, you notice that there's no major difference when committing changes to a SQL database in Microsoft Fabric compared to other code repositories.

- **Commit to source control:** You can commit database objects to source control, converting the live database into code. This process reads object definitions from the database and writes them to the repository.
- **Update from source control:** You can update database objects from the contents of source control. The code is validated before applying a differential change to the database.
- **History tracking:** Users can view the history of database objects in the source control system, providing a clear record of changes and facilitating collaboration.

Explore performance capabilities

SQL Database in Fabric offers intelligent performance capabilities like monitoring, and automatic index creation and tuning.


Monitor with performance dashboard

The Performance Dashboard in Fabric SQL Database simplifies the user experience by removing the complexities of monitoring and operation. It allows users to fully use the SQL database engine's capabilities, catering to various workloads in Fabric.

The dashboard offers different levels of metrics visibility to accommodate users with varying SQL expertise. Beginners can access basic query performance metrics, while intermediate and advanced users can view more detailed information.

You can access the performance dashboard by right-clicking on the context button (the three dots) in the item view, then select **Open performance summary**.

Alternatively, you can access the home toolbar in the **Query Editor** window, and select **Performance summary**.



The performance dashboard helps users understand their database performance and receive alerts for any issues. It's designed to assist application developers in detecting and resolving performance bottlenecks early, ensuring an intuitive and efficient user experience.

Explore automatic tuning

Automatic tuning is a built-in capability that applies machine learning to optimize your query performance. It automatically identifies tuning opportunities and implements them to enhance your database's efficiency.

In SQL database in Fabric, indexes are managed dynamically, with a graph showing the count of created, dropped, and reverted indexes over time, and a table listing the indexes created by the tool, including details like schema name, table name, index name, status, key columns, included columns, and creation and drop dates.

You can monitor automatic indexing on the **Automatic index** tab in the performance dashboard.

3. Secure databases in Fabric

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/3-secure-database>

Secure databases in Fabric

Completed

In a database, users have access to several capabilities aimed at safeguarding sensitive information. These security measures are capable of securing or masking data from users or roles without proper authorization, ensuring data protection across both the database and the SQL analytics endpoints. This ensures a smooth and secure user experience, with no need for changes to the existing applications.

Understand security capabilities

Users who are often well-versed in the SQL engine and adept at using T-SQL, find databases in Microsoft Fabric easy to use.

This is because SQL Database in Microsoft Fabric is powered by the same SQL engine they're familiar with, enabling them to perform complex queries and data manipulations. The SQL engine's wide range of security features further allows for sophisticated security mechanism at the database level.

Workspaces roles

Designed to provide different levels of access and control within the workspace. You can assign users to the various workspace roles such as **Admin**, **Member**, **Contributor**, and **Viewer**. These roles are crucial for maintaining the security and efficiency of SQL database operations within an organization.

Item permissions

Individual databases can have item permissions assigned directly to them. The main intent behind assigning such permissions is to facilitate the sharing of the SQL database for downstream use. You can review permissions for a SQL database, its SQL analytics endpoint, or its default semantic model by navigating to the item in the workspace and selecting the **Manage permissions** quick action.

Note


Granting item permissions has no impact on the security metadata inside the database.

Data protection

You can use T-SQL to grant specific permissions to users for more precise control. SQL database supports a range of data protection features that enable administrators to shield sensitive data from unauthorized access.

It includes object-level security for database objects, column-level security for table columns, row-level security for table rows using `WHERE` clause filters, and dynamic data masking to obscure sensitive data like email addresses.

You can manage database-level roles from Fabric portal by selecting **Security** and then choosing **Manage SQL security**.



4. Explore Copilot for SQL Database

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/4-facilitate-database-tasks-with-copilot>

Explore Copilot for SQL Database

Completed

In today's fast-paced tech world, understanding AI is essential for enhancing applications and staying competitive. SQL Database in Fabric is key to this transformation, offering a robust platform for integrating AI. With capabilities like Natural Language to SQL, code completion, quick actions, and intelligent insights, it empowers users to harness AI, and boost application performance.

These tools enable the creation of intelligent, responsive, and efficient applications that meet modern user demands.

Use Copilot for SQL database in Fabric

Microsoft Copilot is integrated with SQL database in Fabric, enhancing SQL management and troubleshooting. It boosts productivity by offering natural language to SQL conversion and self-help for users.

Copilot can automatically correct T-SQL code errors as they occur. By sharing context with the active query tab, Copilot offers helpful suggestions to fix SQL query errors seamlessly. You can also generate T-SQL queries by asking questions in natural language.

With the **natural language to SQL** capability, which translates natural language queries into SQL within the query editor in Fabric portal, database interactions become more intuitive. This integration allows Copilot to answer questions based on your database context, as Copilot is schema-aware.

Which employees have completed more than two projects this quarter?
List all products that were out of stock last month.
Provide the rank of each department by employee satisfaction score and show the department
Create a primary key constraint named PK_Product_ProductID to the column ProductID in the t

Copilot in Fabric SQL database includes several key features:

- **Chat pane:** Use natural language to ask questions and generate T-SQL queries
- **Code completion:** Automatic suggestions while typing T-SQL code
- **Quick actions:** Explain and Fix buttons in the query editor toolbar to help understand and correct queries

To access Copilot features in the Fabric portal, open your SQL database and select the **Copilot** button in the ribbon to open the chat pane, or use the **Explain** and **Fix** quick action buttons in the query editor toolbar.

Note

Copilot for SQL database does not use data in tables to generate T-SQL suggestions.

5. Simplify data management with mirroring and virtualization

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/5-simplify-data-management>


Simplify data management with mirroring and virtualization

Completed

Organizations often face challenges in managing and analyzing their data due to the complexity of integrating various data sources, ensuring data consistency, and maintaining real-time data availability. SQL Database in Microsoft Fabric addresses these challenges by providing a unified platform that simplifies data integration, enhances data consistency, and ensures near real-time data availability.

Integrate with mirroring

One of the key features of SQL Database in Microsoft Fabric is its ability to mirror databases from Azure SQL Database directly into Fabric's OneLake. Additionally, SQL Database in Fabric is automatically mirrored for analytics purposes, with data continuously replicated into OneLake in near real-time.



This mirroring process ensures that data is continuously replicated in near real-time, eliminating the need for complex Extract, Transform, Load (ETL) processes. By doing so, it reduces the total cost of ownership and accelerates the time-to-insight, allowing businesses to unlock business intelligence, artificial intelligence, data engineering, data science, and data sharing scenarios.

After you initiate a mirroring process, you can monitor the replication status by selecting the **Monitor replication** option from the **Replication** tab. If there are no updates in the source tables, the engine will back off and resume regular polling after detecting updated data.

Database name
AdventureWorksLT
Status
Running
Database ID
aaaaaaaa-0000-1111-2222-bbbbbbbbbbbb

Table list

Name	Status	Rows replicated	Last completed
[dbo].[BuildVersion]	Running	0	--
[SalesLT].[ProductDescription]	Running	0	--
[SalesLT].[ProductModelProdu...	Running	0	--
[SalesLT].[Product]	Running	0	--
[SalesLT].[Address]	Running	0	--

To learn more about how to configure mirrored databases, see [Tutorial: Configure Microsoft Fabric mirrored databases from Azure SQL Database](#).

Explore data virtualization

Data virtualization in SQL Database in Fabric is a capability that allows you to access and manipulate data from various sources without the need to physically move or copy the data. This approach provides a unified view of data, enabling seamless integration and analysis across different platforms.

These features enable scenarios such as querying Parquet, CSV, and Delta tables available in a lakehouse.

Capability	Definition	Query Example
Database scoped credential	Allows you to create credentials that can be used to access external data sources securely.	<pre>CREATE DATABASE SCOPED CREDENTIAL MyCredential WITH IDENTITY = 'USER IDENTITY';</pre>
External data source	This enables you to define external data sources, such as files stored in OneLake.	<pre>'abfss://aaaaaaa-0000-1111-2222-bbbbbbbbbbbb@<onelake_account_name>.dfs.fabric.microsoft.com/bbbbbbbb-1111-2222-3333-cccccccccc/Files/parquet/data1.parquet';</pre>
External file format	This capability lets you specify the format of external files, such as Parquet, CSV, and Delta files.	<pre>CREATE EXTERNAL FILE FORMAT MyFileFormat WITH (FORMAT_TYPE = DELIMITEDTEXT, FORMAT_OPTIONS (FIELD_TERMINATOR = ',', STRING_DELIMITER = '"'));</pre>
External table	This allows you to create tables that reference data stored outside the SQL database.	<pre>CREATE EXTERNAL TABLE MyExternalTable (Column1 INT, Column2 NVARCHAR(50)) WITH (LOCATION = 'myfolder/myfile.csv', DATA_SOURCE = MyExternalDataSource, FILE_FORMAT = MyFileFormat);</pre>

6. Exercise: Work with SQL Database in Microsoft Fabric

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/6-exercise-query-data-sql-database>

Exercise: Work with SQL Database in Microsoft Fabric

Completed

Now it's your chance to work with SQL queries, manage data, and analyze patterns in a SQL database in Microsoft Fabric.

In this exercise, you create a database in Fabric and load it with sample data. You then query the database and combine data from external data sources. Lastly, you create a database role and secure your data.

Note

To complete this exercise, you need an [Azure subscription](#).

Launch the exercise and follow the instructions.

[Launch Exercise](#)

7. Module assessment

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/7-knowledge-check>

Module assessment

Completed

8. Summary

<https://learn.microsoft.com/en-us/training/modules/get-started-sql-database-microsoft-fabric/8-summary>

Summary

Completed

In this module, you have learned about SQL Database in Microsoft Fabric, a Software as a Service (SaaS) transactional database built on Azure SQL Database. The module covered the creation and management of operational databases within the Fabric environment, automatic data replication into OneLake, and conversion to Parquet in near real-time. You also learned about the importance of source control in managing SQL databases in Microsoft Fabric, the platform's intelligent performance capabilities, and the various security measures provided to protect sensitive information. The module also touched on the role of workspace roles and item permissions in controlling access and facilitating sharing.

The main takeaways from this module include understanding how to create a new SQL database in Fabric, query a SQL database using similar tools available for Azure SQL Database, and manage features like object-level, column-level, and row-level security. You also learned about the integration of AI in SQL Database in Fabric, with features like Natural Language to SQL, code completion, quick actions, and intelligent insights. The module also highlighted the role of Microsoft Copilot in enhancing SQL management and troubleshooting. Lastly, you learned about data management simplification through mirroring and data virtualization, and how to secure your data by creating a database role.

Additional reading

- [Microsoft Fabric Documentation](#)
- [SQL database in Microsoft Fabric](#)
- [Overview of Copilot in Fabric](#)
- [What is the admin portal?](#)
- [Share your SQL database and manage permissions](#)
- [What is Mirroring in Fabric?](#)