

Secure data access in Microsoft Fabric

1. Introduction

<https://learn.microsoft.com/en-us/training/modules/secure-data-access-in-fabric/1-introduction>

Introduction

Completed

Security in Microsoft Fabric is optimized for securing data for specific use cases. Different users need the ability to perform various actions in Fabric to fulfill their job responsibilities. Fabric facilitates this by allowing you to grant users access to specific data workloads through workspace and item permissions, compute permissions, and OneLake data access roles (preview).

Secure data by use case

A security use case in Fabric refers to a set of users needing data access and accessing data in a specific way. Once a use case is identified, Fabric permissions associated with that use case can be configured. Suppose you work at a healthcare company with multiple systems that store data, such as electronic health records (EHR), insurance claims data, clinical trial data, patient and disease registries, and administrative data. Different users within your company or partner organizations need to view, transform, analyze, aggregate, and use this data to derive business insights. Users need access to different Fabric compute engines, items, and workspaces to perform their jobs effectively:

- **Data engineers** need access to data in a lakehouse to develop downstream data products.
- **Business or data analysts** need to query data to answer business questions.
- **Data scientists** need to access data in a lakehouse and consume it through Apache Spark to create models and experiments.
- **Report creators** need to build reports to share with report consumers.
- **Report consumers** need to view data in Power BI reports to make decisions.

In this module, you'll learn how to use the Fabric access control features available to secure your data and provide your team with the necessary access within Fabric to perform their job duties. You'll explore Fabric's multi-layer security model and how to use it to manage data access.

By the end of this module, you'll know how to configure security for an entire workspace, for individual Fabric data items, and how to apply granular permissions within Fabric data items.

2. Understand the Fabric security model

<https://learn.microsoft.com/en-us/training/modules/secure-data-access-in-fabric/2-understand-fabric-security-model>

Understand the Fabric security model

Completed

Data access in organizations is often restricted by users' responsibilities, and roles and by an organization's Fabric deployment patterns, and data architecture. Fabric has a flexible, multi-layer security model that allows you to configure security to accommodate different data access requirements. Having the ability to control permissions at different layers means you can adhere to the principle of least privilege, restricting user permissions to only what's needed to perform job tasks.


Fabric has three security levels and they're evaluated sequentially to determine whether a user has data access. The order of evaluation for access is:

1. Microsoft Entra ID authentication: checks if the user can authenticate to the Azure identity and access management service, Microsoft Entra ID.
2. Fabric access: checks if the user can access Fabric.
3. Data security: checks if the user can perform the action they've requested on a table or file.

The third level, data security, has several building blocks that can be configured individually or together to align with different access requirements. The primary access controls in Fabric are:

- Workspace roles
- Item permissions
- Compute or granular permissions
- OneLake data access controls (preview)

It's helpful to envision these building blocks in a hierarchy to understand how access controls can be applied individually or together.



A *workspace* in Fabric enables you to distribute ownership and access policies using *workspace roles*. Within a workspace, you can create Fabric data *items* like lakehouses, data warehouses, and semantic models. *Item permissions* can be inherited from a workspace role or set individually by sharing an item. When workspace roles provide too much access, items can be shared using item permissions to ensure proper security.

Within each data item, *granular engine permissions* such as Read, ReadData, or ReadAll can be applied.

Compute or granular permissions can be applied within a specific compute engine in Fabric, like the SQL Endpoint or semantic model.

Fabric data items store their data in OneLake. Access to data in the lakehouse can be restricted to specific files or folders using the role-based-access control (RBAC) feature called OneLake data access controls (preview).

3. Configure workspace and item permissions

<https://learn.microsoft.com/en-us/training/modules/secure-data-access-in-fabric/3-configure-workspace-and-item-permissions>

Configure workspace and item permissions

Completed

Workspaces are environments where users can collaborate to create groups of items. Items are the resources you can work with in Fabric such as lakehouses, warehouses, and reports. Workspace roles are preconfigured sets of permissions that let you manage what users can do and access in a Fabric workspace.

Item permissions control access to individual Fabric items within a workspace. Item permissions let you either adjust the permissions set by a workspace role or give a user access to one or more items within a workspace without adding the user to a workspace role.

Let's consider some scenarios where you would need to configure data access using workspace roles and item permissions.

Understand workspace roles

Suppose you work at a health care company as the Fabric security admin. You need to set up access for a new data engineer. The data engineer needs the ability to:

- Create Fabric items in an existing workspace
- Read all data in an existing lakehouse that's in the same workspace where they can create Fabric items

Workspace roles control what users can do and access within a Fabric workspace. There are four workspace roles and they apply to *all items* within a workspace. Workspace roles can be assigned to individuals, security groups, Microsoft 365 groups, and distribution lists. Users can be assigned to the following roles:

- **Admin** - Can view, modify, share, and manage all content and data in the workspace, and manage permissions.
- **Member** - Can view, modify, and share all content and data in the workspace.
- **Contributor** - Can view and modify all content and data in the workspace.
- **Viewer** - Can view all content and data in the workspace, but can't modify it.


Tip

For a full list of the permissions associated with workspace roles, see: [Roles in workspaces](#)

To meet the access requirements for the new data engineer, you can assign them the workspace **Contributor** role. This gives them access to modify content in the workspace, including creating Fabric items like lakehouses. The contributor role would also allow them to read data in the existing lakehouse.

Assign workspace roles

Users can be added to workspace roles from the **Manage access** button from within a workspace. Add a user by entering the user's name and selecting the workspace role to assign them in the **Add people** dialogue.




Configure item permissions

Item permissions control access to individual Fabric items within a workspace. Item permission can be used to give a user access to one or more items within a workspace without adding the user to a workspace role or can be used with workspace roles.

Suppose that after a few months of having **Contributor** access on a workspace, a data engineer no longer needs to create Fabric items and now only needs to view a single lakehouse and read data in it.

Since the engineer no longer needs to view all items in the workspace, the **Contributor** workspace role can be removed and item permissions on the lakehouse can be configured so the engineer will only be able to see the lakehouse metadata and data and nothing else in the workspace. This item access configuration helps you adhere to the principle of least privilege, where the engineer only has access to what's needed to perform their job duties.

An item can be shared and item permissions can be configured by selecting on the ellipsis (...) next to a Fabric item in a workspace and then selecting **Manage permissions**.



In the **Grant people access** window that appears after selecting **Manage permissions**, if you add the user and don't select any of the checkboxes under **Additional permissions**, the user will have read access to the lakehouse metadata. The user won't have access to the underlying data in the lakehouse. To grant the engineer the ability to read data and not just metadata, **Read all SQL endpoint data** or **Read all Apache Spark** can be selected.

Grant people access



samplelakehouse

People you share this Lakehouse with can open it and its SQL endpoint and read the default dataset. To allow them to read directly in the Lakehouse, grant additional permissions.

A

Alice

×

Additional permissions

- ☒ Read all SQL endpoint data ⓘ
- ☒ Read all Apache Spark ⓘ
- ☐ Build reports on the default semantic model

Notification Options

- ☒ Notify recipients by email

Add a message (optional)

 Depending on which additional permissions you select, recipients will have different access to the SQL endpoint, default dataset, and data in the lakehouse. For details, view lakehouse permissions documentation.

Grant

Back

Tip

Each Fabric data item has its own security model. To learn more about permissions that can be granted when a lakehouse or other Fabric data item is shared see:

- [Warehouse](#)
- [Lakehouse](#)
- [Semantic model](#)

4. Apply granular permissions

<https://learn.microsoft.com/en-us/training/modules/secure-data-access-in-fabric/4-apply-granular-permissions>

Apply granular permissions

Completed

When the permissions provided by workspace roles or item permissions are insufficient, granular permissions like table and row-level security and file and folder access can be set through the:

- SQL analytics endpoint
- OneLake data access roles (preview)
- Warehouse
- Semantic model

Configure data access through the SQL analytics endpoint in a lakehouse


Data in a lakehouse can be read through the SQL analytics endpoint. Each Lakehouse has an autogenerated SQL analytics endpoint that can be used to transition between the *lake* view of the lakehouse and the *SQL* view of the lakehouse. The *lake* view supports data engineering and Apache Spark and the *SQL* view of the same lakehouse allows you to create views, functions, stored procedures and to apply SQL security and object level permissions.

Data in a Fabric lakehouse is stored with the following folder structure:

- /Files
- /Tables

View the SQL analytics endpoint view of the lakehouse

The SQL analytics endpoint is used to read data in the /Tables folder of the lakehouse using T-SQL.



Apply granular permissions to the lakehouse using T-SQL

Using the SQL analytics endpoint, granular T-SQL permissions can be applied to SQL objects using Data Control Language (DCL) commands such as:


- [GRANT](#)
- [DENY](#)
- [REVOKE](#)

Row-level security, column-level security, and dynamic data masking can also be applied using the SQL analytics endpoint. See:

- [Row-level security](#)
- [Column-level security](#)
- [Dynamic data masking](#)

Configure data access through the lake view of the lakehouse

The lake view of the lakehouse is used to read data in the /Tables and /Files folder of the lakehouse.




Use OneLake data access roles to secure data

Workspace and item permissions provide coarse access to data in a lakehouse. To further refine data access, folders in the lake view of the lakehouse can be secured using OneLake data access roles

(preview). You can create custom roles within a lakehouse and grant read permissions only to specific folders in OneLake. Folder security is inheritable to all subfolders. To create a custom OneLake data access role:

1. Select **Manage OneLake data access (preview)** from the menu in the lake view of the lakehouse.




2. In the **New Role** window, create a new role name and select the folders to grant access to.
3. Once the role is created, assign a user or group to the role and select the permissions to assign.

Tip

For more information on how OneLake RBAC permissions are evaluated with workspace and item permissions, see: [How OneLake RBAC permissions are evaluated with Fabric permissions](#)

Configure granular warehouse permissions

Granular permissions can be applied to warehouses using the SQL analytics endpoint, similar to the way the endpoint is used for the lakehouse. The same permissions can be applied: GRANT, REVOKE, and DENY and row-level security, column-level security, and dynamic data masking.



Configure Semantic model permissions

A user's role in a workspace implicitly grants them permission on the semantic models in a workspace. Semantic models allow for security to be defined using DAX. More granular permission can be applied using row-level security (RLS). To learn more about the managing RLS or permissions on the semantic model see:

- [Semantic model permissions](#)
- [Row-level security \(RLS\) with Power BI](#)

5. Exercise: Secure data access in Microsoft Fabric

<https://learn.microsoft.com/en-us/training/modules/secure-data-access-in-fabric/5-exercise-secure-data-access>

Exercise: Secure data access in Microsoft Fabric

Completed

Now it's your chance to secure data access in Microsoft Fabric.

In this exercise, you learn how to secure data access in Fabric using the concepts explored in this module.

Note

- You need a Microsoft Fabric trial license with the Fabric preview enabled in your tenant. See [Getting started with Fabric](#) to enable your Fabric trial license.
- To complete the exercises in this lab, you'll need two users: one user should be assigned the Workspace Admin role, and the other will be assigned permissions throughout this lab. To assign roles to workspaces see [Give access to your workspace](#).

Launch the exercise and follow the instructions.

[Launch Exercise](#)

6. Module assessment

Module assessment

Completed

7. Summary

<https://learn.microsoft.com/en-us/training/modules/secure-data-access-in-fabric/7-summary>

Summary

Completed

In this module, you learned how to use the access control features available across various Fabric engines to secure your data and provide your team with the necessary access within Fabric to perform their job duties. You explored Fabric's multi-layer security model and how to use it to manage data access.

The main takeaways from this module include understanding how to configure security for an entire workspace, and for individual Fabric data items, and how to apply granular permissions.

For more reading, you can refer to the following URLs:

- [Microsoft Fabric security white paper](#)
- [Microsoft Fabric permission model](#)
- [Build common data architectures with OneLake in Microsoft Fabric](#)
- [How to secure data for common Fabric data architectures](#)
- [End-to-end security scenario](#)