

Azure DP-700 Microsoft Fabric Data Engineer Associate

Free Practice Questions Preview

Here are 35 sample questions to help you get started. Unlock the full exam to access all 60+ questions with detailed explanations.

Question 1 : Implement and manage an analytics solution

Contoso Retail has separate dev, test, and prod Fabric workspaces for a lakehouse solution. The team wants controlled promotion of notebooks, pipelines, and lakehouse items with approvals and minimal manual copying. What should you configure?

- A. A OneLake shortcut from dev to prod so all items are shared automatically
- B. A Dataflow Gen2 refresh schedule in each workspace
- C. Workspace Viewer access for data scientists in all environments
- D. Deployment pipelines for the workspaces and Git integration for source-controlled items**

Answer: D

Deployment pipelines promote Fabric content across stages while Git tracks source-controlled changes. A shortcut shares data paths, not application lifecycle promotion or approvals.

Question 2 : Implement and manage an analytics solution

City Power Utility needs engineers to collaborate on Fabric notebooks using pull requests and branches. They also need to revert a bad change. Which capability best supports this requirement?

- A. Configure Git integration for the workspace and connect it to a repository branch**
- B. Export every notebook manually before each release
- C. Use sensitivity labels on notebooks
- D. Create a semantic model refresh schedule

Answer: A

Git integration provides branching, history, and revert workflows. Manual export is error-prone and does not provide collaborative version control.

Question 3 : Implement and manage an analytics solution

Tails핀 Toys stores customer profiles in a Fabric lakehouse. Analysts should query only approved tables, while engineers can update notebooks and pipelines. In this case, what is the best governance pattern?

- A. Give all analysts the workspace Admin role
- B. Combine workspace roles with item-level permissions and granular data permissions where required**
- C. Put all data in one folder and rely only on naming conventions

D. Disable OneLake access for the workspace

Answer: B

Fabric governance is layered: workspace roles control collaboration, item permissions control access, and granular permissions restrict data. Admin access gives excessive privileges.

Question 4 : Implement and manage an analytics solution

Northwind Health must prevent BI developers from seeing a sensitive column in a warehouse table while allowing normal query access to the rest of the table. Which feature should be used?\n

A. A workspace deployment pipeline

B. A Spark session setting

C. Column-level security or dynamic data masking based on the access requirement

D. A OneLake shortcut to a different workspace

Answer: C

Column-level security or masking protects sensitive columns during query access. Deployment pipelines move artifacts and do not enforce column visibility.

Question 5 : Implement and manage an analytics solution

Northwind Health wants users to discover trusted datasets and lakehouse items in Fabric. The data owner has validated the content and wants to mark it as reliable. What should be applied?\n

A. Increase the Spark driver size

B. Create a pipeline parameter

C. Use KQL update policy on the table

D. Endorse the item as promoted or certified according to the governance process

Answer: D

Endorsement communicates that content is trusted or certified for reuse. Spark sizing and pipeline parameters do not signal trust.

Question 6 : Implement and manage an analytics solution

City Power Utility must track who accessed or changed Fabric items that contain sales orders. Which capability is most appropriate?\n

A. Microsoft Fabric audit logs

B. Dataflow Gen2 computed tables only

C. Spark broadcast joins

D. Warehouse result-set caching

Answer: A

Audit logs are used to review user and activity events for governance and compliance. Computed tables transform data but do not provide audit history.

Question 7 : Implement and manage an analytics solution

Datum Insurance needs a low-code transformation that business data owners can maintain. The source is SQL Server, and the result must land in OneLake. Which Fabric item is usually the best fit?\n

- A. A KQL database function only
- B. Dataflow Gen2**
- C. A deployment pipeline
- D. A workspace domain

Answer: B

Dataflow Gen2 is the low-code data preparation option for ingesting and transforming data into Fabric destinations. KQL functions are not the typical low-code ingestion choice.

Question 8 : Implement and manage an analytics solution

Northwind Health has complex Python logic for parsing support tickets, using custom libraries and several joins. The process must be parameterized and called by a pipeline. Which item should implement the transformation?\n

- A. A sensitivity label
- B. A workspace role assignment
- C. A Fabric notebook invoked from a pipeline with parameters**
- D. A manual CSV upload

Answer: C

Notebooks are appropriate for custom Python or PySpark transformations and can be orchestrated with pipeline parameters. Labels and role assignments do not transform data.

Question 9 : Implement and manage an analytics solution

Blue Yonder Logistics is implementing governance and lifecycle management for Fabric items that process customer profiles. The architecture review rejected manual steps and over-permissive access. Which two actions should you choose? Each correct answer presents part of the solution.\n

- A. Give all developers Admin access in production and copy notebooks manually
- B. Use workspace domains only and skip item-level permissions
- C. Store production logic in personal workspaces and export files before release
- D. Connect the development workspace to Git and promote validated content with deployment pipelines**

Answer: D

Git plus deployment pipelines supports controlled lifecycle management. Admin access and manual copying are common traps because they bypass governance and repeatability.

Question 10 : Implement and manage an analytics solution

Litware Manufacturing needs a pipeline to pass the processing date and source folder dynamically to a notebook. Which design is best?\n

A. Use pipeline parameters and dynamic expressions, then pass them to the notebook activity

- B. Hard-code the date in the notebook
- C. Create separate workspaces for each date
- D. Use only workspace Viewer permissions

Answer: A

Pipeline parameters and dynamic expressions support reusable orchestration. Hard-coding dates creates brittle runs and is a common exam trap.

Question 11 : Implement and manage an analytics solution

City Power Utility wants to centralize ownership of workspaces by business area and make Fabric assets easier to discover by domain. Which setting should be configured?\n

- A. Spark shuffle partitions
- B. Fabric domain workspace settings**
- C. T-SQL statistics
- D. Eventstream capture settings

Answer: B

Domain workspace settings organize workspaces by business domains for discoverability and governance. Spark and query settings do not classify ownership domains.

Question 12 : Implement and manage an analytics solution

Contoso Retail has a Fabric workspace where Spark jobs often fail because the default compute is too small. What should an administrator tune first?\n

- A. Item endorsement
- B. A warehouse primary key
- C. Spark workspace settings such as pool/session configuration**
- D. The audit log retention label

Answer: C

Spark workspace settings control default Spark behavior and compute configuration. Endorsement affects trust metadata, not job capacity.

Question 13 : Implement and manage an analytics solution

Blue Yonder Logistics wants users to access customer profiles through OneLake while enforcing data access consistently across engines where possible. Which area should be configured?

- A. Only a Power BI theme
- B. Only notebook comments
- C. Only a deployment pipeline description
- D. OneLake security, combined with item and data permissions**

Answer: D

OneLake security helps govern data access across OneLake-backed items. Themes and comments do not enforce access.

Question 14 : Implement and manage an analytics solution

Alpine Ski House must separate who can manage a workspace from who can read a specific warehouse item. Which statement is correct?

- A. Workspace roles and item permissions are separate layers and should both be considered**
- B. A user with Viewer always becomes warehouse administrator
- C. Sensitivity labels replace all access control
- D. Git branches determine production permissions

Answer: A

Fabric uses multiple security layers. Workspace roles do not remove the need to configure item and data permissions.

Question 15 : Implement and manage an analytics solution

Blue Yonder Logistics has a raw-to-curated process for sales orders. The team needs sequential steps, retry handling, and dependency management. Which Fabric item is most appropriate for orchestration?

- A. A sensitivity label
- B. A data pipeline**
- C. A workspace description
- D. A SQL endpoint shortcut only

Answer: B

Data pipelines orchestrate activities, dependencies, retries, schedules, and parameters. Labels and descriptions do not

execute process flows.

Question 16 : Implement and manage an analytics solution

Fabrikam Finance is deciding between Dataflow Gen2, a pipeline, and a notebook. The requirement is to copy files, call a notebook, and send a notification after completion. What should be the controlling item?\n

- A. A Dataflow Gen2 item only
- B. A sensitivity label
- C. A pipeline that orchestrates copy, notebook, and notification activities**
- D. A KQL queryset only

Answer: C

Pipelines are designed to coordinate multiple activities. Dataflow Gen2 is mainly for low-code data prep, not end-to-end orchestration.

Question 17 : Implement and manage an analytics solution

Contoso Retail wants to apply the same schema changes to a Fabric warehouse across environments through a source-controlled database project. What should the team use?\n

- A. Manual table recreation in production
- B. Eventstream windows
- C. Lakehouse table vacuum only
- D. Fabric database projects with lifecycle management**

Answer: D

Database projects support source-controlled database schema lifecycle. Manual recreation is risky and not a controlled deployment approach.

Question 18 : Implement and manage an analytics solution

Alpine Ski House is implementing governance and lifecycle management for Fabric items that process IoT telemetry. The workload has a strict SLA and must be observable. Which two actions should you choose? Each correct answer presents part of the solution.\n

- A. Give all developers Admin access in production and copy notebooks manually
- B. Use workspace domains only and skip item-level permissions
- C. Connect the development workspace to Git and promote validated content with deployment pipelines**
- D. Store production logic in personal workspaces and export files before release

Answer: C

Git plus deployment pipelines supports controlled lifecycle management. Admin access and manual copying are

common traps because they bypass governance and repeatability.

Question 19 : Implement and manage an analytics solution

Datum Insurance classifies sales orders as confidential and wants the classification to follow Fabric items and appear to users. In this case, what should be used?

- A. Spark repartition
- B. Sensitivity labels**
- C. Pipeline concurrency
- D. KQL ingestion batching

Answer: B

Sensitivity labels classify and communicate data sensitivity. Repartitioning and batching are performance techniques, not classification controls.

Question 20 : Implement and manage an analytics solution

A release from test to production at Contoso Retail must include only validated Fabric items and avoid changing unrelated artifacts. What is the best practice?

- A. Give every developer production Admin access
- B. Copy files directly in OneLake Explorer
- C. Use deployment pipeline stages and compare/select items before deployment**
- D. Disable Git before deployment

Answer: C

Deployment pipelines allow controlled comparison and deployment between stages. Granting broad production Admin access is a governance risk.

Question 21 : Implement and manage an analytics solution

City Power Utility wants a repeatable nightly process that loads curated tables after all raw files arrive. Failures should be visible in Fabric monitoring. Which combination is best?

- A. Workspace color theme and item endorsement
- B. Manual notebook execution by an analyst
- C. Only a SQL view without orchestration
- D. Pipeline schedule with activities for ingestion, transformation, and monitoring**

Answer: D

A scheduled pipeline provides repeatable orchestration and observable activity runs. Manual runs are not reliable for production.

Question 22 : Implement and manage an analytics solution

Contoso Retail has a table containing energy meter readings. Different business units must see only their own rows when querying. Which control directly addresses this?\n

- A. Row-level security or equivalent granular data access control
- B. Spark executor memory
- C. Git pull request policy only
- D. Deployment pipeline stage naming

Answer: A

Row-level security restricts visible rows based on user context. Compute and naming conventions do not filter query results.

Question 23 : Implement and manage an analytics solution

City Power Utility is planning Fabric governance. Which pair of actions most directly supports trustworthy discovery and compliance review?\n

- A. Increase Spark partitions and cache every table
- B. Apply sensitivity labels and use audit logs
- C. Disable workspaces and use personal exports
- D. Use only CSV files and avoid item permissions

Answer: B

Labels help classify data, and audit logs support compliance review. Spark tuning helps performance, not governance evidence.

Question 24 : Implement and manage an analytics solution

Blue Yonder Logistics wants to manage Dataflows Gen2 defaults and staging behavior for a workspace. Where should this be configured?\n

- A. A KQL query window
- B. The warehouse statistics page
- C. Dataflows Gen2 workspace settings
- D. A semantic model relationship view

Answer: C

Dataflows Gen2 workspace settings control workspace-level Dataflow behavior. KQL and warehouse tools do not configure Dataflow workspace defaults.

Question 25 : Implement and manage an analytics solution

Contoso Retail needs to ensure production items are changed only through a controlled release path. What is the strongest pattern?\n

- A. Let developers edit production notebooks directly
- B. Use only file names to indicate versions
- C. Use a single workspace for dev, test, and prod
- D. Use Git for development history and deployment pipelines for environment promotion**

Answer: D

Git plus deployment pipelines supports controlled lifecycle management. Direct production edits and one shared workspace increase release risk.

Question 26 : Implement and manage an analytics solution

Contoso Retail has separate dev, test, and prod Fabric workspaces for a lakehouse solution. The team wants controlled promotion of notebooks, pipelines, and lakehouse items with approvals and minimal manual copying. Given least privilege, automation, and production reliability, what should you configure?\n

- A. Deployment pipelines for the workspaces and Git integration for source-controlled items**
- B. A OneLake shortcut from dev to prod so all items are shared automatically
- C. A Dataflow Gen2 refresh schedule in each workspace
- D. Workspace Viewer access for analytics engineers in all environments

Answer: A

Deployment pipelines promote Fabric content across stages while Git tracks source-controlled changes. A shortcut shares data paths, not application lifecycle promotion or approvals.

Question 27 : Implement and manage an analytics solution

Datum Insurance is implementing governance and lifecycle management for Fabric items that process customer profiles. The platform team wants an approach that works across dev, test, and production. Which two actions should you choose? Each correct answer presents part of the solution.\n

- A. Give all developers Admin access in production and copy notebooks manually
- B. Connect the development workspace to Git and promote validated content with deployment pipelines**
- C. Use workspace domains only and skip item-level permissions
- D. Store production logic in personal workspaces and export files before release

Answer: B

Git plus deployment pipelines supports controlled lifecycle management. Admin access and manual copying are common traps because they bypass governance and repeatability.

Question 28 : Implement and manage an analytics solution

Datum Insurance stores clickstream events in a Fabric lakehouse. Analysts should query only approved tables, while engineers can update notebooks and pipelines. What is the best governance pattern?\n

- A. Give all analysts the workspace Admin role
- B. Put all data in one folder and rely only on naming conventions
- C. Combine workspace roles with item-level permissions and granular data permissions where required**
- D. Disable OneLake access for the workspace

Answer: C

Fabric governance is layered: workspace roles control collaboration, item permissions control access, and granular permissions restrict data. Admin access gives excessive privileges.

Question 29 : Implement and manage an analytics solution

Adventure Works must prevent external partners from seeing a sensitive column in a warehouse table while allowing normal query access to the rest of the table. Which feature should be used?\n

- A. A workspace deployment pipeline
- B. A Spark session setting
- C. A OneLake shortcut to a different workspace
- D. Column-level security or dynamic data masking based on the access requirement**

Answer: D

Column-level security or masking protects sensitive columns during query access. Deployment pipelines move artifacts and do not enforce column visibility.

Question 30 : Implement and manage an analytics solution

Adventure Works wants users to discover trusted datasets and lakehouse items in Fabric. The data owner has validated the content and wants to mark it as reliable. Given least privilege, automation, and production reliability, what should be applied?\n

- A. Endorse the item as promoted or certified according to the governance process**
- B. Increase the Spark driver size
- C. Create a pipeline parameter
- D. Use KQL update policy on the table

Answer: A

Endorsement communicates that content is trusted or certified for reuse. Spark sizing and pipeline parameters do not signal trust.

Question 31 : Implement and manage an analytics solution

Blue Yonder Logistics must track who accessed or changed Fabric items that contain payment events. Which capability is most appropriate?\n

- A. Dataflow Gen2 computed tables only
- B. Microsoft Fabric audit logs**
- C. Spark broadcast joins
- D. Warehouse result-set caching

Answer: B

Audit logs are used to review user and activity events for governance and compliance. Computed tables transform data but do not provide audit history.

Question 32 : Implement and manage an analytics solution

Proseware Media needs a low-code transformation that business data owners can maintain. The source is REST APIs, and the result must land in OneLake. Which Fabric item is usually the best fit?\n

- A. A KQL database function only
- B. A deployment pipeline
- C. Dataflow Gen2**
- D. A workspace domain

Answer: C

Dataflow Gen2 is the low-code data preparation option for ingesting and transforming data into Fabric destinations. KQL functions are not the typical low-code ingestion choice.

Question 33 : Implement and manage an analytics solution

Proseware Media has complex Python logic for parsing claims transactions, using custom libraries and several joins. The process must be parameterized and called by a pipeline. Which item should implement the transformation?\n

- A. A sensitivity label
- B. A workspace role assignment
- C. A manual CSV upload
- D. A Fabric notebook invoked from a pipeline with parameters**

Answer: D

Notebooks are appropriate for custom Python or PySpark transformations and can be orchestrated with pipeline parameters. Labels and role assignments do not transform data.

Question 34 : Implement and manage an analytics solution

Proseware Media wants a pipeline to run when a file arrives in a monitored storage location, rather than only at 02:00 each day. Given least privilege, automation, and production reliability, what should be configured?

- A. An event-based trigger for the pipeline
- B. A higher Spark executor count
- C. A certified endorsement
- D. A warehouse view

Answer: A

Event-based triggers start orchestration when an event occurs, such as file arrival. Increasing Spark executors affects compute, not trigger behavior.

Question 35 : Implement and manage an analytics solution

Alpine Ski House needs a pipeline to pass the processing date and source folder dynamically to a notebook. Which design is best?

- A. Hard-code the date in the notebook
- B. Use pipeline parameters and dynamic expressions, then pass them to the notebook activity
- C. Create separate workspaces for each date
- D. Use only workspace Viewer permissions

Answer: B

Pipeline parameters and dynamic expressions support reusable orchestration. Hard-coding dates creates brittle runs and is a common exam trap.

Unlock All 60+ Questions

Get the complete Q&A package with detailed explanations, topic analytics, and exam-accurate practice.

From €29.00

Visit: <https://cert-pass.com/exams/azure-dp-700-microsoft-fabric-data-engineer-associate>

CERT-PASS

cert-pass.com

© 2026 Cert-Pass. This material is for personal use only. Do not distribute.