

Fabric Certification Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

Copyright © 2025 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Questions

SAMPLE

- 1. What does the AS CLONE OF statement do in a warehouse?**
 - A. Creates an independent copy of the table's data**
 - B. Creates a replica of the table by copying metadata only**
 - C. Deep copies both data and metadata of the table**
 - D. Not applicable in warehouse environments**
- 2. Which tool should the Research division use to group their workspaces?**
 - A. The Fabric Admin Portal**
 - B. The Azure Portal**
 - C. The Microsoft Entra admin center**
 - D. The Azure Logic Apps**
- 3. What should you select to enable partition discovery in the Copy activity?**
 - A. Enable staging**
 - B. Enable partition discovery**
 - C. Set Mode to Append**
 - D. Select the partition column**
- 4. What is the first modification needed to ensure read-write access to a dataset via XMLA endpoint?**
 - A. Modify the DS1 settings.**
 - B. Modify the WS1 settings.**
 - C. Modify the C1 settings.**
 - D. Modify the Tenant1 settings.**
- 5. Which three steps are necessary to hide a column in Power BI for a specific role?**
 - A. Open the model in Tabular Editor, Select the column, Set Object Level Security to Default.**
 - B. Open the model in Power BI Desktop, Select the column, Set Object Level Security to None.**
 - C. Open the model in Tabular Editor, Select the column, Set Object Level Security to None.**
 - D. Open the model in Power BI Desktop, Select the column, Set the Hidden property to true.**

- 6. What is the correct command to ensure a Delta table reads all existing rows in addition to newly added ones?**
- A. write**
 - B. update**
 - C. append**
 - D. add**
- 7. When visualizing customer metrics, which tool provides insights into data distributions and outliers?**
- A. A. Column charts**
 - B. B. Heat maps**
 - C. C. Pie charts**
 - D. D. Scatter plots**
- 8. When reading a CSV file and saving as a Delta table, which option helps in optimizing the read operation?**
- A. Inferschema**
 - B. Partitioning**
 - C. Removing unnecessary columns**
 - D. Explicitly defining the schema**
- 9. To identify minimum and maximum values in a dataflow, which of the following is NOT directly selected in Data view options?**
- A. A. Show column value distribution**
 - B. B. Enable column profile**
 - C. C. Show column profile in details pane**
 - D. D. Show custom visualizations**
- 10. If you want to connect to a managed table in Lakehouse1, which action is allowed?**
- A. Read the table.**
 - B. Update the table.**
 - C. Delete the table.**
 - D. Drop the table.**

Answers

SAMPLE

1. B
2. A
3. B
4. C
5. C
6. C
7. A
8. C
9. D
10. A

SAMPLE

Explanations

SAMPLE

1. What does the AS CLONE OF statement do in a warehouse?

- A. Creates an independent copy of the table's data**
- B. Creates a replica of the table by copying metadata only**
- C. Deep copies both data and metadata of the table**
- D. Not applicable in warehouse environments**

The AS CLONE OF statement creates a replica of a table by copying its metadata only. This means that while the structure of the original table — including its definitions and constraints — is duplicated, the actual data within the table itself is not copied over. This functionality is particularly useful in environments where you want to maintain a lightweight version or reference of a table without the overhead of duplicating the data. In data warehousing, this capability allows analysts and developers to create test scenarios or reporting contexts without the need to create large volumes of redundant data, ensuring both efficiency and resource management. The metadata allows for the new table to operate with a defined structure, enabling further actions like querying or altering columns without the data itself being present.

2. Which tool should the Research division use to group their workspaces?

- A. The Fabric Admin Portal**
- B. The Azure Portal**
- C. The Microsoft Entra admin center**
- D. The Azure Logic Apps**

The Fabric Admin Portal is the appropriate tool for the Research division to group their workspaces because it is specifically designed to manage and organize resources within the Fabric, facilitating efficient workspace management. This portal provides features that allow users to create, configure, and manage workspaces, making it ideal for grouping various research projects and their corresponding workspaces under a cohesive administration interface. Using the Fabric Admin Portal helps ensure that all team members in the Research division have access to the organized structure of their workspaces, enabling easy collaboration and resource management. This centralized management capability is crucial for tracking projects, sharing insights, and maintaining version control across the workspaces. Other tools mentioned, such as the Azure Portal and the Microsoft Entra admin center, serve different purposes related to broader Azure resource management and identity governance, while Azure Logic Apps focus on automating workflows. These tools, although powerful in their respective functions, do not provide the specialized capabilities needed to group and manage Fabric workspaces effectively.

3. What should you select to enable partition discovery in the Copy activity?

- A. Enable staging**
- B. Enable partition discovery**
- C. Set Mode to Append**
- D. Select the partition column**

To enable partition discovery in the Copy activity, selecting the option to enable partition discovery is the essential step. This functionality allows the data pipeline to automatically identify and utilize partitions in data sources that support partitioning. By enabling this feature, the data integration service can dynamically adjust to variations in data segmentation, allowing for more efficient data processing and loading. Partition discovery optimizes how data is copied by enabling the activity to recognize and access specific subsets of data based on defined partitions, such as those defined by date or any other criterion. This greatly enhances performance by minimizing the amount of data that needs to be processed or transferred at one time, thus improving efficiency and reducing resource consumption. Other options may involve aspects of the data management process, like staging data or configuring the mode and partition column, but they do not specifically address the mechanism needed to discover partitions in the data source itself. Activating partition discovery is the crucial action that ensures the Copy activity can effectively leverage existing data partitions.

4. What is the first modification needed to ensure read-write access to a dataset via XMLA endpoint?

- A. Modify the DS1 settings.**
- B. Modify the WS1 settings.**
- C. Modify the C1 settings.**
- D. Modify the Tenant1 settings.**

To ensure read-write access to a dataset via the XMLA endpoint, modifying the settings of the specific workspace or dataset that you want to grant access to is crucial. In this case, the option that involves modifying the settings related to the dataset (denoted as C1) indicates that these settings directly affect how data can be accessed and manipulated through the XMLA. Without making the appropriate changes within the specific configuration that governs how the dataset interacts with the XMLA interface, users would not be able to perform write operations on it. These settings typically include permissions and access levels that explicitly enable read and write functionalities. Therefore, focusing on the correct dataset settings is essential for achieving the desired access. Modifications to other settings such as those for specific services or broader tenant settings may not directly relate to enabling access for a particular dataset, which is why they do not address the core issue at hand regarding XMLA endpoint access.

5. Which three steps are necessary to hide a column in Power BI for a specific role?

- A. Open the model in Tabular Editor, Select the column, Set Object Level Security to Default.**
- B. Open the model in Power BI Desktop, Select the column, Set Object Level Security to None.**
- C. Open the model in Tabular Editor, Select the column, Set Object Level Security to None.**
- D. Open the model in Power BI Desktop, Select the column, Set the Hidden property to true.**

To hide a column in Power BI for a specific role, you need to ensure that you are working with appropriate features that manage visibility and security at the model level. The correct procedure involves using Tabular Editor, where you can manage advanced features such as Object Level Security. When you select the column in Tabular Editor and set the Object Level Security to None, it means that you are removing any restrictions related to that column for users assigned to that specific role. This effectively hides the column from the view of users with that role, aligning with the goal of controlling data visibility based on role assignments. Other options may mention using Power BI Desktop or improperly configuring Object Level Security settings, which do not achieve the desired outcome of hiding the column appropriately for specific user roles. Tabular Editor is essential in managing these settings effectively and is the recommended tool for more advanced data modeling tasks.

6. What is the correct command to ensure a Delta table reads all existing rows in addition to newly added ones?

- A. write**
- B. update**
- C. append**
- D. add**

The command that ensures a Delta table reads all existing rows as well as newly added ones is the append command. When using this command, the system is instructed to add new data to the existing data set while retaining all previously stored records. In Delta tables, which are designed for handling large volumes of data efficiently, it is crucial to maintain all existing rows when incorporating new data. The append action is specifically designed for this purpose—adding new records without overriding or excluding any of the existing information. Thus, when a Delta table is queried after an append operation, users will see the complete dataset, inclusive of both historical and newly inserted records.

7. When visualizing customer metrics, which tool provides insights into data distributions and outliers?

A. A. Column charts

B. B. Heat maps

C. C. Pie charts

D. D. Scatter plots

The best tool for visualizing customer metrics in a way that provides insights into data distributions and outliers is the scatter plot. Scatter plots are particularly valuable for showing the relationship between two quantitative variables, allowing you to see patterns, trends, and anomalies in the data. In a scatter plot, each point represents an observation, and by examining the distribution of these points, you can quickly identify clusters, gaps, and outliers. This makes it easier to monitor how various customer metrics relate to one another at different levels and understand the underlying structure of your data. Other visualization types such as column charts, heat maps, and pie charts serve different purposes. Column charts are useful for comparing categories, heat maps provide insights into the density of data points based on color gradients, and pie charts are effective for displaying relative proportions of a whole. However, none of these options effectively highlight distributions and outliers as clearly as scatter plots do. This makes scatter plots the ideal tool for the given task.

8. When reading a CSV file and saving as a Delta table, which option helps in optimizing the read operation?

A. Inferschema

B. Partitioning

C. Removing unnecessary columns

D. Explicitly defining the schema

When saving a CSV file as a Delta table, optimizing read operations is crucial for performance, and one effective strategy is to remove unnecessary columns. This practice streamlines the data stored in the Delta table, which can reduce the amount of data that needs to be read during queries. By eliminating columns that are not required for specific analyses or operations, you minimize the data footprint, leading to faster read times. This also benefits storage efficiency, as less data means a reduced amount of disk space is used. Consequently, when queries are executed, the system can process the smaller dataset more swiftly, improving overall performance. While inferring the schema, partitioning, and explicitly defining the schema can all have benefits in different contexts, they do not directly contribute to reducing the amount of data processed in the same way that removing unnecessary columns does. Hence, focusing on the relevant data ensures that the read operation is optimized for speed and efficiency.

9. To identify minimum and maximum values in a dataflow, which of the following is NOT directly selected in Data view options?

- A. A. Show column value distribution**
- B. B. Enable column profile**
- C. C. Show column profile in details pane**
- D. D. Show custom visualizations**

The process of identifying minimum and maximum values in a dataflow primarily relies on tools that analyze and profile data attributes. The options focusing on column value distribution, enabling column profiles, and displaying detailed profiles are all integral to assessing the characteristics of the data, including its range and distribution. The option concerned with showing custom visualizations, however, does not directly aid in identifying minimum and maximum values. While custom visualizations can enhance understanding of data patterns and insights, they do not specifically facilitate the quantification or extraction of extreme values directly from the data columns. Thus, this option does not contribute to the immediate task of determining the min-max range, making it the correct answer in this context. The capabilities associated with profiling and distribution specifically target statistical analysis, which is essential when looking for minimum and maximum values, whereas custom visualizations are often used for broader interpretative analytics.

10. If you want to connect to a managed table in Lakehouse1, which action is allowed?

- A. Read the table.**
- B. Update the table.**
- C. Delete the table.**
- D. Drop the table.**

Connecting to a managed table in Lakehouse1 typically permits a user to perform operations like reading the data stored within that table. A managed table in this context refers to a table where the framework manages both the schema and the data. This allows users to easily access and query the data without needing to handle the underlying storage details. Reading the table is fundamental for data analysis and reporting, which likely aligns with the expected use cases in a Lakehouse environment. Operations like updating, deleting, or dropping the table can require additional permissions or may alter the structure and content of the table, which goes beyond merely accessing the data. Hence, reading is the most basic and common operation that users can perform without needing elevated privileges or risking data integrity.