

Hot DP-600 Questions - DP-600 Free Learning Cram



BONUS!!! Download part of ActualTorrent DP-600 dumps for free: https://drive.google.com/open?id=1H_QlaferLRwcZ6w-3zLt-4AkwibGmlj

Our DP-600 exam dumps are required because people want to get succeed in IT field by clearing the certification exam. Passing DP-600 practice exam is not so easy and need to spend much time to prepare the training materials, that's the reason that so many people need professional advice for DP-600 Exam Prep. The DP-600 dumps pdf are the best guide for them passing test.

Microsoft DP-600 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">Implement and manage semantic models: This section of the exam measures the skills of architects and focuses on designing and optimizing semantic models to support enterprise-scale analytics. It evaluates understanding of storage modes and implementing star schemas and complex relationships, such as bridge tables and many-to-many joins. Architects must write DAX-based calculations using variables, iterators, and filtering techniques. The use of calculation groups, dynamic format strings, and field parameters is included. The section also includes configuring large semantic models and designing composite models. For optimization, candidates are expected to improve report visual and DAX performance, configure Direct Lake behaviors, and implement incremental refresh strategies effectively.
Topic 2	<ul style="list-style-type: none">Prepare data: This section of the exam measures the skills of engineers and covers essential data preparation tasks. It includes establishing data connections and discovering sources through tools like the OneLake data hub and the real-time hub. Candidates must demonstrate knowledge of selecting the appropriate storage type—lakehouse, warehouse, or eventhouse—depending on the use case. It also includes implementing OneLake integrations with Eventhouse and semantic models. The transformation part involves creating views, stored procedures, and functions, as well as enriching, merging, denormalizing, and aggregating data. Engineers are also expected to handle data quality issues like duplicates, missing values, and nulls, along with converting data types and filtering. Furthermore, querying and analyzing data using tools like SQL, KQL, and the Visual Query Editor is tested in this domain.
Topic 3	<ul style="list-style-type: none">Maintain a data analytics solution: This section of the exam measures the skills of administrators and covers tasks related to enforcing security and managing the Power BI environment. It involves setting up access controls at both workspace and item levels, ensuring appropriate permissions for users and groups. Row-level, column-level, object-level, and file-level access controls are also included, alongside the application of sensitivity labels to classify data securely. This section also tests the ability to endorse Power BI items for organizational use and oversee the complete development lifecycle of analytics assets by configuring version control, managing Power BI Desktop projects, setting up deployment pipelines, assessing downstream impacts from various data assets, and handling semantic model deployments using XMLA endpoint. Reusable asset management is also a part of this domain.

Top Hot DP-600 Questions | High Pass-Rate Microsoft DP-600: Implementing Analytics Solutions Using Microsoft Fabric 100% Pass

Our DP-600 exam materials are the most reliable products for customers. If you need to prepare an exam, we hope that you can choose our DP-600 study guide as your top choice. In the past ten years, we have overcome many difficulties and never give up. And we have quickly grown up as the most influential company in the market. And our DP-600 preparation questions are the most popular among the candidates.

Microsoft Implementing Analytics Solutions Using Microsoft Fabric Sample Questions (Q96-Q101):

NEW QUESTION # 96

You have a Fabric workspace named Workspace1.

Workspace1 contains multiple semantic models, including a model named Model1. Model1 is updated by using an XMLA endpoint. You need to increase the speed of the write operations of the XMLA endpoint.

What should you do?

- A. Configure Model1 to use the Direct Lake storage format.
- B. Delete any unused columns from Model1.
- C. Delete any unused semantic models from Workspace1.
- D. **Select Large semantic model storage format for Workspace1.**

Answer: D

Explanation:

When using XMLA endpoints to manage and update semantic models in Microsoft Fabric, the performance of write operations (such as processing, structural changes, or metadata deployments from Tabular Editor) is directly influenced by the storage format and how the model is persisted.

Why Option A is Correct

- * By default, Fabric semantic models use the Small semantic model storage format.
- * To improve write operations performance through XMLA, you must change the workspace setting to use the Large semantic model storage format.
- * The large format uses more efficient storage techniques, supports partitioning, and handles larger models with optimized write capabilities.
- * This setting is applied at the workspace level and impacts all semantic models within that workspace, including Model1.

This is explicitly documented in Microsoft's guidance: Large semantic model storage format is required when using XMLA write operations for large or frequently updated models.

Why the Other Options Are Incorrect

B). Configure Model1 to use the Direct Lake storage format.

- * Direct Lake mode is designed for query performance (reading data directly from OneLake in delta format without import/duplication).
- * It improves query latency and freshness but does not improve XMLA write operations, which deal with model metadata and structural updates.

C). Delete any unused semantic models from Workspace1.

- * Deleting unused semantic models helps manage capacity and storage but does not increase the speed of XMLA endpoint write operations.
- * Workspace storage overhead does not directly impact the write throughput of XMLA operations.

D). Delete any unused columns from Model1.

- * Removing unused columns reduces the memory footprint and can improve query performance.
- * However, it does not directly improve the speed of XMLA write operations. The bottleneck in XMLA writes is tied to the storage format, not the model size alone.

Summary

To increase the speed of XMLA write operations on semantic models, you must enable the Large semantic model storage format at the workspace level. This setting ensures better handling of writes and metadata operations via the XMLA endpoint.

References

- * Large models in Power BI and Microsoft Fabric
- * Use the XMLA endpoint in Microsoft Fabric

* Manage capacities in Microsoft Fabric

NEW QUESTION # 97

Case Study 2 - Litware, Inc

Overview

Litware, Inc. is a manufacturing company that has offices throughout North America. The analytics team at Litware contains data engineers, analytics engineers, data analysts, and data scientists.

Existing Environment

Fabric Environment

Litware has been using a Microsoft Power BI tenant for three years. Litware has NOT enabled any Fabric capacities and features.

Available Data

Litware has data that must be analyzed as shown in the following table.

Description	Original source	Total size
Customer data	Customer relationship management (CRM) system	50 MB
Product data	Customer relationship management (CRM) system	200 MB
Customer satisfaction surveys	SurveyMonkey	500 GB

The Product data contains a single table and the following columns.

Name	Data type
ProductID	Integer
ProductName	String
ProductCategory	String
ListPrice	Decimal

The customer satisfaction data contains the following tables:

- Survey
- Question
- Response

For each survey submitted, the following occurs:

- One row is added to the Survey table.
- One row is added to the Response table for each question in the survey.
- The Question table contains the text of each survey question. The third question in each survey response is an overall satisfaction score. Customers can submit a survey after each purchase.

User Problems

The analytics team has large volumes of data, some of which is semi-structured. The team wants to use Fabric to create a new data store.

Product data is often classified into three pricing groups: high, medium, and low. This logic is implemented in several databases and semantic models, but the logic does NOT always match across implementations.

Requirements

Planned Changes

Litware plans to enable Fabric features in the existing tenant. The analytics team will create a new data store as a proof of concept (PoC). The remaining Litware users will only get access to the Fabric features once the PoC is complete. The PoC will be completed by using a Fabric trial capacity. The following three workspaces will be created:

- AnalyticsPOC: Will contain the data store, semantic models, reports pipelines, dataflow, and notebooks used to populate the data store
- DataEngPOC: Will contain all the pipelines, dataflows, and notebooks used to populate OneLake
- DataSciPOC: Will contain all the notebooks and reports created by the data scientists. The following will be created in the AnalyticsPOC workspace:
 - A data store (type to be decided)
 - A custom semantic model
 - A default semantic model

Interactive reports

The data engineers will create data pipelines to load data to OneLake either hourly or daily depending on the data source. The analytics engineers will create processes to ingest, transform, and load the data to the data store in the AnalyticsPOC workspace daily. Whenever possible, the data engineers will use low-code tools for data ingestion. The choice of which data cleansing and transformation tools to use will be at the data engineers' discretion.

All the semantic models and reports in the Analytics POC workspace will use the data store as the sole data source.

Technical Requirements

The data store must support the following:

- Read access by using T-SQL or Python
- Semi-structured and unstructured data
- Row-level security (RLS) for users executing T-SQL queries

Files loaded by the data engineers to OneLake will be stored in the Parquet format and will meet Delta Lake specifications.

Data will be loaded without transformation in one area of the AnalyticsPOC data store. The data will then be cleansed, merged, and transformed into a dimensional model. The data load process must ensure that the raw and cleansed data is updated completely before populating the dimensional model. The dimensional model must contain a date dimension. There is no existing data source for the date dimension. The Litware fiscal year matches the calendar year. The date dimension must always contain dates from 2010 through the end of the current year.

The product pricing group logic must be maintained by the analytics engineers in a single location. The pricing group data must be made available in the data store for T-SQL queries and in the default semantic model. The following logic must be used:

- List prices that are less than or equal to 50 are in the low pricing group.
- List prices that are greater than 50 and less than or equal to 1,000 are in the medium pricing group.
- List prices that are greater than 1,000 are in the high pricing group.

Security Requirements

Only Fabric administrators and the analytics team must be able to see the Fabric items created as part of the PoC.

Litware identifies the following security requirements for the Fabric items in the AnalyticsPOC workspace:

- Fabric administrators will be the workspace administrators.
- The data engineers must be able to read from and write to the data store. No access must be granted to datasets or reports.
- The analytics engineers must be able to read from, write to, and create schemas in the data store. They also must be able to create and share semantic models with the data analysts and view and modify all reports in the workspace.
- The data scientists must be able to read from the data store, but not write to it. They will access the data by using a Spark notebook
- The data analysts must have read access to only the dimensional model objects in the data store. They also must have access to create Power BI reports by using the semantic models created by the analytics engineers.
- The date dimension must be available to all users of the data store.
- The principle of least privilege must be followed.

Both the default and custom semantic models must include only tables or views from the dimensional model in the data store. Litware already has the following Microsoft Entra security groups:

FabricAdmins: Fabric administrators

- AnalyticsTeam: All the members of the analytics team
- DataAnalysts: The data analysts on the analytics team
- DataScientists: The data scientists on the analytics team
- DataEngineers: The data engineers on the analytics team
- AnalyticsEngineers: The analytics engineers on the analytics team

Report Requirements

The data analysts must create a customer satisfaction report that meets the following requirements:

- Enables a user to select a product to filter customer survey responses to only those who have purchased that product.
- Displays the average overall satisfaction score of all the surveys submitted during the last 12 months up to a selected date.
- Shows data as soon as the data is updated in the data store.
- Ensures that the report and the semantic model only contain data from the current and previous year.
- Ensures that the report respects any table-level security specified in the source data store.
- Minimizes the execution time of report queries.

You need to recommend a solution to prepare the tenant for the PoC.

Which two actions should you recommend performing from the Fabric Admin portal? Each correct answer presents part of the solution.

NOTE: Each correct answer is worth one point.

- A. Enable the Allow Azure Active Directory guest users to access Microsoft Fabric option for specific security groups.
- B. **Enable the Users can create Fabric items option for specific security groups.**
- C. Enable the Users can try Microsoft Fabric paid features option for the entire organization.
- D. Enable the Users can create Fabric items option and exclude specific security groups.
- E. **Enable the Users can try Microsoft Fabric paid features option for specific security groups.**

Answer: B,E

NEW QUESTION # 98

Hotspot Question

You have a Fabric workspace that contains a warehouse named Warehouse1. Warehouse1 contains the following data.

Table name	Column name	Data type
Employee	EmployeeID	Int
Employee	EmployeeName	Varchar(128)
Employee	EmployeePosition	Varchar(64)
Contract	EmployeeID	Int
Contract	ContractType	Varchar(64)
Contract	StartDate	Datetime2
Contract	EndDate	Datetime2

You need to create a T-SQL statement that will denormalize the tables and include the ContractType and StartDate attributes in the results. The solution must meet the following requirements:

- Include attributes from matching rows in the Contract table.
- Ensure that all the rows from the Employee table are preserved.
- Return the total number of employees per contract type for all the contract types that have more than two employees.

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```

WITH result AS(
    SELECT
        e.EmployeeID
        , e.EmployeeName
        , e.EmployeePosition
        , c.ContractType
    FROM Employee AS e
    CROSS JOIN Contract AS c ON c.EmployeeID = e.EmployeeID)
    SELECT
        COUNT(DISTINCT EmployeeID) AS TotalEmployees, ContractType
    FROM result
    GROUP BY ContractType
    HAVING COUNT(DISTINCT EmployeeID) > 2

```

Answer:

Explanation:

Answer Area

```
WITH result AS(
    SELECT
        e.EmployeeID
        , e.EmployeeName
        , e.EmployeePosition
        , c.ContractType
    FROM Employee AS e
    LEFT OUTER JOIN
        Contract AS c
        ON c.EmployeeID = e.EmployeeID
)
SELECT
    COUNT(DISTINCT EmployeeID) AS TotalEmployees, ContractType
FROM result
GROUP BY ContractType
HAVING
    COUNT(DISTINCT EmployeeID) > 2
```



NEW QUESTION # 99

You are developing a Microsoft Power BI semantic model. Two tables in the data model are not connected in a physical relationship. You need to establish a virtual relationship between the tables. Which DAX function should you use?

- A. CROSSFILTER()
- B. TREATAS()
- C. USERELATIONSHIP()
- D. PATH()

Answer: B

NEW QUESTION # 100

You have a Fabric tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a subfolder named Subfolder1 that contains CSV files.

You need to convert the CSV files into the delta format that has V-Order optimization enabled. What should you do from Lakehouse explorer?

- A. Create a new shortcut in the Files section.
- B. Create a new shortcut in the Tables section.
- C. Use the Load to Tables feature.
- D. Use the Optimize feature.

Answer: C

Explanation:

With "Load to tables" : tables are always loaded using the Delta Lake table format with V-Order optimization enabled.

<https://learn.microsoft.com/en-us/fabric/data-engineering/load-to-tables/load-to-table-capabilities-overview>

NEW QUESTION # 101

• • • • •

Whereas the other two ActualTorrent DP-600 exam questions formats are concerned, both are customizable practice tests, provide real time environment, track your progress, and help you overcome mistakes. The desktop Microsoft DP-600 Practice Test software is compatible with Windows computers. The web based practice exam is supported by all browsers and operating systems.

DP-600 Free Learning Cram <https://www.actualtorrent.com/DP-600-questions-answers.html>

P.S. Free & New DP-600 dumps are available on Google Drive shared by ActualITorrent: https://drive.google.com/open?id=1H_QlaferLRwcZ6w-3zLt-4AkwiibGmIj