

DP-800 Free Exam - Braindump DP-800 Pdf



We have authoritative production team made up by thousands of experts helping you get hang of our Developing AI-Enabled Database Solutions study question and enjoy the high quality study experience. We will update the content of DP-800 test guide from time to time according to recent changes of examination outline and current policies, so that every examiner can be well-focused and complete the exam focus in the shortest time. Besides, our DP-800 Exam Questions can help you optimize your learning method by simplifying obscure concepts so that you can master better. One more to mention, with our DP-800 test guide, there is no doubt that you can cut down your preparing time in 20-30 hours of practice before you take the exam.

Our DP-800 certification files are the representative masterpiece and leading in the quality, service and innovation. We collect the most important information about the test DP-800 certification and supplement new knowledge points which are produced and compiled by our senior industry experts and authorized lecturers and authors. We provide the auxiliary functions such as the function to stimulate the real exam to help the clients learn our DP-800 Quiz materials efficiently and pass the DP-800 exam.

>> DP-800 Free Exam <<

Microsoft DP-800 Exam | DP-800 Free Exam - 100% Pass Rate Offer of Braindump DP-800 Pdf

Have you learned Prep4King Microsoft DP-800 exam dumps? Why do the people that have used Prep4King dumps sing its praises? Do you really want to try it whether it have that so effective? Hurry to click Prep4King.com to download our certification training materials. Every question provides you with demo and if you think our exam dumps are good, you can immediately purchase

it. After you purchase DP-800 Exam Dumps, you will get a year free updates. Within a year, only if you would like to update the materials you have, you will get the newer version. With the dumps, you can pass Microsoft DP-800 test with ease and get the certificate.

Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> Implement AI capabilities in database solutions: This domain covers designing and managing external AI models and embeddings, implementing full-text, semantic vector, and hybrid search strategies, and building retrieval-augmented generation (RAG) solutions that connect database outputs with language models.
Topic 2	<ul style="list-style-type: none"> Secure, optimize, and deploy database solutions: This domain focuses on implementing data security measures like encryption, masking, and row-level security, optimizing query performance, managing CI/CD pipelines using SQL Database Projects, and integrating SQL solutions with Azure services including Data API builder and monitoring tools.
Topic 3	<ul style="list-style-type: none"> Design and develop database solutions: This domain covers designing and building database objects such as tables, views, functions, stored procedures, and triggers, along with writing advanced T-SQL code and leveraging AI-assisted tools like GitHub Copilot and MCP for SQL development.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q18-Q23):

NEW QUESTION # 18

You have an Azure SQL database named AdventureWorksDB that contains a table named dbo.Employee.

You have a C# Azure Functions app that uses an HTTP-triggered function with an Azure SQL input binding to query dbo.Employee.

You are adding a second function that will react to row changes in dbo.Employee and write structured logs.

You need to configure AdventureWorksDB and the app to meet the following requirements:

- * Changes to dbo.Employee must trigger the new function within five seconds.
- * Each invocation must process no more than 100 changes.

Which two database configurations should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Enable change tracking at the database level.
- **B. Set Sql_Trigger_PollingIntervalMs to 5000.**
- C. Create an AFTER trigger on dbo.Employee for Data Manipulation Language (DML).
- D. Enable change data capture (CDC) for dbo.Employee table changes
- **E. Enable change tracking on the dbo. Employee table.**
- F. SetSql Trigger MaxBatchSize to 100.

Answer: B,E

Explanation:

Azure Functions' Azure SQL trigger requires change tracking to be enabled on the source table. Microsoft's SQL trigger documentation states that setting up change tracking for the Azure SQL trigger requires two steps: enable change tracking on the database and enable change tracking on the table being monitored.

Since the question asks specifically which database configurations you should perform, enabling change tracking on dbo.Employee is one of the required database-side steps.

To meet the latency requirement that changes trigger the function within five seconds, the relevant trigger setting is Sql_Trigger_PollingIntervalMs. Microsoft documents this setting as the delay, in milliseconds, between processing each batch of changes, and a value of 5000 means the trigger polls every 5 seconds.

A few clarifications about the other options:

- * B is not the documented setting name. The documented app setting is Sql_Trigger_BatchSize or host setting MaxBatchSize, not "SetSql Trigger MaxBatchSize". The screenshot wording suggests a distractor.
- * D is also required in practice for the trigger to work, but the question asks for two answers and includes the polling setting plus the table-level CT setting as the actionable choices presented.
- * F is wrong because the Azure SQL trigger uses change tracking, not CDC.

NEW QUESTION # 19

You have a Microsoft SQL Server 2025 instance that has a managed identity enabled.

You have a database that contains a table named `dbo.ManualChunks`. `dbo.ManualChunks` contains product manuals.

A retrieval query already returns the top five matching chunks as `nvarchar(max)` text.

You need to call an Azure OpenAI REST endpoint for chat completions. The solution must provide the highest level of security.

You write the following Transact-SQL code.

What should you insert at line 02?

- A.
- B.
- C.
- D.
- E.

Answer: D

Explanation:

The correct answer is Option B because the requirement is to call an Azure OpenAI REST endpoint from SQL Server 2025 while providing the highest level of security, and the instance already has a managed identity enabled. For Microsoft's SQL AI features, the preferred secure pattern is to use a database scoped credential with `IDENTITY = 'Managed Identity'` instead of storing an API key. Microsoft documents that SQL Server 2025 supports managed identity for external AI endpoints, and for Azure OpenAI the credential secret uses the Cognitive Services resource identifier: `{ "resourceid " : "https://cognitiveservices.azure.com" }`.

So line 02 should be:

```
WITH IDENTITY = 'Managed Identity',
```

```
SECRET = '{ "resourceid " : "https://cognitiveservices.azure.com" }';
```

 Why the other options are incorrect:

* A and D use HTTP header or query-string credentials with an API key, which is less secure than managed identity because a secret key must be stored and rotated manually. Microsoft recommends managed identity where supported to avoid embedded secrets.

* C mixes Managed Identity with an api-key secret, which is not the correct pattern for Azure OpenAI managed-identity authentication.

* E uses an invalid identity value for this scenario. The accepted credential identities for external REST endpoint calls include `HTTPHeaderHeaders`, `HTTPHeaderQueryString`, `Managed Identity`, and `Shared Access Signature`.

Because the endpoint is Azure OpenAI and the question explicitly asks for the highest security, managed identity with the Cognitive Services resource ID is the Microsoft-aligned answer.

NEW QUESTION # 20

You have an Azure SQL database that contains the following SQL graph tables:

* A NODE table named `dbo.Person`

* An EDGE table named `dbo.Knows`

Each row in `dbo.Person` contains the following columns:

* `Personid` (int)

* `DisplayName` (nvarchar(100))

You need to use a HATCH operator and exactly two directed `Knows` relationships to return the `Personid` and `DisplayName` of people that are reachable from the person identified by an input parameter named `@startPersonid`.

Which Transact-SQL query should you use?

- A.
- B.
- C.
- D.

Answer: D

Explanation:

The correct query is Option D because it starts from the input person and uses exactly two directed `Knows` edges in a single `MATCH` pattern:

MATCH(p1-(k1)->p2-(k2)->p3)

Microsoft documents that SQL Graph uses the MATCH predicate in the WHERE clause to express graph traversal patterns over node and edge tables, and directed relationships are written with arrow syntax such as node1-(edge)->node2.

Why D is correct:

- * It anchors the starting node with p1.PersonId = @StartPersonId.
- * It traverses two directed hops : p1 -> p2 -> p3.
- * It returns p3.PersonId, p3.DisplayName, which are the people reachable in exactly two Knows relationships.

Why the others are wrong:

- * A filters on DisplayName = DisplayName, which is unrelated to the required input parameter and does not correctly anchor the start node.
- * B reverses the traversal direction in the pattern.
- * C uses two separate MATCH predicates instead of the required single two-hop directed pattern. The proper graph pattern syntax supports chaining the hops directly in one MATCH expression.

Topic 1, Contoso Case Study

Existing Environment

Contoso has an Azure subscription in North Europe that contains the corporate infrastructure. The current infrastructure contains a Microsoft SQL Server 2017 database. The database contains the following tables.

The FeedbackJson column has a full-text index and stores JSON documents in the following format.

The support staff at Contoso never has the unmask permission.

Requirements

Contoso is deploying a new Azure SQL database that will become the authoritative data store for the following:

- * AI workloads
- * Vector search
- * Modernized API access
- * Retrieval Augmented Generation (RAG) pipelines

Sometimes the ingestion pipeline fails due to malformed JSON and duplicate payloads.

The engineers at Contoso report that the following dashboard query runs slowly.

```
SELECT VehicleId, LastUpdatedUtc, EngineStatus, BatteryHealth FROM dbo.VehicleHealthSummary where fleetId = gFleetId  
ORDER BY LastUpdatedUtc DESC;
```

You review the execution plan and discover that the plan shows a clustered index scan. vehicleincidentReports often contains details about the weather, traffic conditions, and location. Analysts report that it is difficult to find similar incidents based on these details.

Planned Changes

Contoso wants to modernize Fleet Intelligence Platform to support AI-powered semantic search over incident reports.

Security Requirements

Contoso identifies the following telemetry requirements:

- * Telemetry data must be stored in a partitioned table.
- * Telemetry data must provide predictable performance for ingestion and retention operations.
- * latitude, longitude, and accuracy JSON properties must be filtered by using an index seek.

Contoso identifies the following maintenance data requirements:

- * Ensure that any changes to a row in the MaintenanceEvents table updates the corresponding value in the LastModifiedDate column to the time of the change.
- * Avoid recursive updates.

AI Search, Embeddings, and Vector indexing

The development team at Contoso will use Microsoft Visual Studio Code and GitHub Copilot and will retrieve live metadata from the databases. Contoso identifies the following requirements for querying data in the FeedbackJson column of the customer-Feedback table:

- * Extract the customer feedback text from the JSON document.
- * Filter rows where the JSON text contains a keyword.
- * Calculate a fuzzy similarity score between the feedback text and a known issue description.
- * Order the results by similarity score, with the highest score first.

NEW QUESTION # 21

Your team is developing an Azure SQL dataset solution from a locally cloned GitHub repository by using Microsoft Visual Studio Code and GitHub Copilot Chat.

You need to disable the GitHub Copilot repository-level instructions for yourself without affecting other users.

What should you do?

- A. Add a - debug flag when you start the GitHub Copilot Chat extension.
- **B. From Visual Studio Code, modify your GitHub Copilot Chat user settings.**

- C. Delete .github/copilot-instructions.md

Answer: B

Explanation:

GitHub documents that repository custom instructions for Copilot Chat can be disabled for your own use in the editor settings, and that doing so does not affect other users. In VS Code, this is controlled through settings related to instruction files, where you can disable the use of repository instruction files for your own environment.

The other options are incorrect:

- * B is not a documented mechanism for disabling repository-level Copilot instructions.
- * C would remove the repository instruction file itself and therefore affect everyone using that repository, which violates the requirement.

NEW QUESTION # 22

You have an Azure SQL database named SalesDB that contains a table named dbo. Articles, dbo.Articles contains two million articles with embeddings. The articles are updated frequently throughout the day.

You query the embeddings by using VECTOR_SEARCH

Users report that semantic search results do NOT reflect the updates until the following day.

You need to ensure that the embeddings are updated whenever the articles change. The solution must minimize CPU usage on SalesDB Which embedding maintenance method should you implement?

- A. Run an hourly Transact-SQL job that regenerates embeddings for all the rows in dbo.Articles.
- B. On dbo.Articles, create a trigger that calls AI GENERATE EMBEDDINGS for each inserted or updated row.
- C. Modify the query to use VECTOR_DISTANCE instead of VECTOR_SEARCH
- D. enable change data capture (CDC) on dbo.Articles and use an Azure Functions app to process CLX changes.

Answer: D

Explanation:

The correct answer is B because the problem is not the vector search operator itself. The problem is that embeddings are becoming stale when article content changes. Microsoft documents that change data capture (CDC) tracks insert, update, and delete operations on source tables, which makes it the right mechanism to identify only the rows that changed.

This also best satisfies the requirement to minimize CPU usage on SalesDB. With CDC, the database only records the row changes, and the embedding regeneration work can be moved to an external process such as an Azure Functions app. That avoids running embedding generation inline inside the database for every update and avoids repeatedly recalculating embeddings for unchanged rows. In contrast, an hourly full-table regeneration would be extremely wasteful on a table with two million frequently updated articles, and a trigger that calls embedding generation per row would push expensive AI work into the transactional path of the database.

Option A is incorrect because changing from VECTOR_SEARCH to VECTOR_DISTANCE does not regenerate embeddings; it only changes the retrieval method. Microsoft states that VECTOR_SEARCH is the ANN search function, while VECTOR_DISTANCE performs exact distance calculation, so neither option addresses stale embedding data.

So the right design is:


- * use CDC to detect only changed articles,
- * process those changes outside the database,
- * regenerate embeddings only for changed rows,
- * write back the refreshed embeddings for current semantic search results.

NEW QUESTION # 23

.....

You may urgently need to attend DP-800 certificate exam and get the certificate to prove you are qualified for the job in some area. But why DP-800 certificate is valuable and useful and can help you a lot? Because passing the test certification can help you prove that you are competent in some area and if you buy our DP-800 Study Materials you will pass the test almost without any problems. We are professional in these career for more than ten years and can give you promised success.

Braindump DP-800 Pdf: <https://www.prep4king.com/DP-800-exam-prep-material.html>

- Latest DP-800 Exam Forum DP-800 Reliable Test Materials  Exam Dumps DP-800 Provider Open www.pdfdumps.com and search for DP-800 to download exam materials for free DP-800 Exam Online

