

Excellent Microsoft Valid Test Syllabus–100% Pass-Rate DP-800 Valuable Feedback



GetValidTest DP-800 Desktop Practice Exam Software: In the Desktop DP-800 practice exam software version of DP-800 practice test is updated and real. The software is useable on Windows-based computers and laptops. There is a demo of the Developing AI-Enabled Database Solutions (DP-800) practice exam which is totally free. Microsoft DP-800 practice test is very customizable and you can adjust its time and number of questions.

The Microsoft DP-800 exam questions on the platform have been gathered by subject matter experts to ensure that they accurately reflect the format and difficulty level of the actual Microsoft DP-800 exam. This makes these Developing AI-Enabled Database Solutions PDF Questions ideal for individuals looking to pass the Microsoft DP-800 Exam on their first try. You can evaluate the product with a free DP-800 demo.

>> Valid DP-800 Test Syllabus <<

DP-800 Valuable Feedback & DP-800 Pass4sure

Our DP-800 exam torrent will not only help you clear exam in your first try, but also enable you prepare exam with less time and effort. There are DP-800 free download trials for your reference before you buy and you can check the accuracy of our questions and answers. Try to Practice DP-800 Exam Pdf with our test engine and you will get used to the atmosphere of the formal test easily.

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 CICD 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 (Q47-Q52):

NEW QUESTION # 47

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. enable change data capture (CDC) on dbo.Articles and use an Azure Functions app to process CLX changes.
- C. Modify the query to use VECTOR_DISTANCE instead of VECTOR_SEARCH.
- D. On dbo.Articles, create a trigger that calls AI_GENERATE_EMBEDDINGS for each inserted or updated row.

Answer: B

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 # 48

You are creating a table that will store customer profiles.

You have the following Transact-SQL code.

□

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Explanation:

* The schema meets the security requirements for PII data. # Yes

* Administrators of the Azure SQL server can see all the rows in dbo.CustomerProfiles when they use an application. # No

* The masking rules will apply even when row-level security (RLS) filters out rows. # No The first statement is Yes because the design combines two relevant SQL security controls for personally identifiable information: Dynamic Data Masking (DDM) on sensitive columns such as FullName, EmailAddress, and PhoneNumber, and Row-Level Security (RLS) to restrict which rows a user can access based on RegionCode. Microsoft documents that DDM limits sensitive data exposure for nonprivileged users, while RLS restricts row access according to the user executing the query. Together, these are valid and appropriate controls for protecting PII in Azure SQL Database.

The second statement is No. Administrative users can view unmasked data because administrative roles effectively have CONTROL, which includes UNMASK. However, that does not mean they automatically see all rows through the application query path defined by the RLS policy. The security policy filters rows based on SUSER_SNAME() and matching RegionCode, so row visibility is governed by the predicate unless the policy is altered or bypassed administratively. DDM and RLS solve different problems: DDM affects how returned values are shown, while RLS affects which rows are returned at all.

The third statement is No because masking only applies to data that is actually returned in the query result set.

Microsoft describes DDM as hiding sensitive data in the result set of a query. If RLS filters a row out, that row is not returned, so there is nothing left for masking to act on. In other words, RLS eliminates inaccessible rows first from the user's perspective, and DDM masks sensitive column values only on rows the user is allowed to see.

NEW QUESTION # 49

You have an Azure SQL database that contains a table named dbo.Products. dbo.Products contains three columns named Embedding, Category, and Price. The Embedding column is defined as VECTOR(1536).

You use AI_GENERATE_EMBEDDINGS and VECTOR_SEARCH to support semantic search and apply additional filters on two columns named Category and Price.

You plan to change the embedding model from text-embedding-ada-002 to text-embedding-3-small. Existing rows already contain embeddings in the Embedding column.

You need to implement the model change. Applications must be able to use VECTOR_SEARCH without runtime errors.

What should you do first?

- A. Convert the Embedding column to nvarchar(max).
- B. Normalize the vector lengths before storing new embeddings.
- **C. Regenerate embeddings for the existing rows.**
- D. Create a vector index on dbo.Products.Embedding.

Answer: C

Explanation:

When you change embedding models, the stored vectors should be treated as belonging to a different embedding space unless you intentionally keep the entire corpus consistent. Microsoft's vector guidance notes that when most or all embeddings are replaced with fresh embeddings from a new model, the recommended practice is to reload the new embeddings and, for large-scale replacement scenarios, consider dropping and recreating the vector index afterward so search quality remains predictable.

This question also says applications must continue to use VECTOR_SEARCH without runtime errors.

VECTOR_SEARCH requires compatible vector dimensions, and the vector column already exists. Azure OpenAI documentation shows that text-embedding-ada-002 is fixed at 1536 dimensions and text-embedding-3-small supports up to 1536 dimensions.

That means the migration can remain compatible with a VECTOR(1536) column, but the right implementation step is still to re-embed the existing rows so the table does not contain a mixed corpus produced by different models.

The other options are not appropriate:

* B normalization does not solve a model migration problem.

* C converting the vector column to nvarchar(max) would break vector-native search design.

* D a vector index improves performance, but it does not migrate old embeddings to the new model.

NEW QUESTION # 50

You need to enable similarity search to provide the analysts with the ability to retrieve the most relevant health summary reports. The solution must minimize latency.

What should you include in the solution?

- A. a standard nonclustered index on the Embeddings (vector (1536)) column
- **B. a vector index on the Embedding* (vector (1536)) column**
- C. a full-text index on the Embeddings (vector (1536)) column
- D. a computed column that manually compares vector values

Answer: B

Explanation:

The correct answer is D because the requirement is to enable similarity search over embedding vectors and to minimize latency . Microsoft documents that CREATE VECTOR INDEX is specifically used to create an index on vector data for approximate nearest neighbor (ANN) search , which is designed to accelerate vector similarity queries compared to exact k-nearest-neighbor scans.

This matches the scenario exactly. The VehicleHealthSummary table already includes an Embeddings (vector (1536)) column. In Microsoft SQL platforms, embeddings are stored in vector columns and queried for semantic similarity. To improve performance and reduce response time, Microsoft recommends a vector index , not a regular B-tree nonclustered index and not a full-text index. A vector index is purpose-built for finding the most similar vectors efficiently.

The other options are not appropriate:

* A would require manual comparison logic and would increase latency rather than minimize it.

* B is incorrect because a standard nonclustered index is not the index type used for vector similarity operations.

* C is incorrect because full-text indexes are for textual token-based search, not numeric vector embeddings.

Microsoft's current documentation is explicit that vector indexes support approximate nearest neighbor search , and that the optimizer can use the ANN index automatically for vector queries. That is the exam- aligned design choice when the goal is fast retrieval of the most relevant health summary reports from an embeddings column.

NEW QUESTION # 51

You have a database named DB1. The schema is stored in a Git repository as an SDK-style SQL database project.

You have a GitHub Actions workflow that already runs dotnet build and produces a database artifact.

You need to add a deployment step that publishes the dacpac file to an Azure SQL database by using the secrets stored in GitHub repository secrets What should you include in the workflow?

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

Answer: D

Explanation:

The correct workflow step is Option C because it uses the Azure SQL GitHub Action to publish a .dacpac file and reads the connection string from GitHub repository secrets , which is exactly what the requirement asks for. Microsoft's Azure SQL GitHub Actions guidance shows using azure/sql-action@v2 with a connection string stored in secrets and a DACPAC path for deployment. The key parts that make C correct are:

* uses: azure/sql-action@v2

* action: publish

* path: bin/Debug/db1.dacpac

* connection-string: \${{ secrets.SQL_CONNECTION_STRING }}

That matches the documented publish pattern for deploying a DACPAC to Azure SQL Database from GitHub Actions. Microsoft and the Azure SQL action documentation both describe Publish as the deployment action for applying a DACPAC to a target database, while Extract is used to create a DACPAC from an existing database, not deploy one.

Why the other options are incorrect:

* A uses an environment variable defined inline with a visible connection string rather than using GitHub repository secrets , which does not meet the requirement.

* B uses action: extract, which would create a DACPAC from a database instead of publishing the existing DACPAC artifact.

* D passes a target connection string to dotnet build, but the question says the workflow already runs dotnet build and produces a database artifact . The missing step is the deployment/publish step, not another build step. Microsoft's SQL project automation guidance separates build the DACPAC from publish the DACPAC .

