Reliable Microsoft DP-800 Exam Book & Latest DP-800 Test Sample

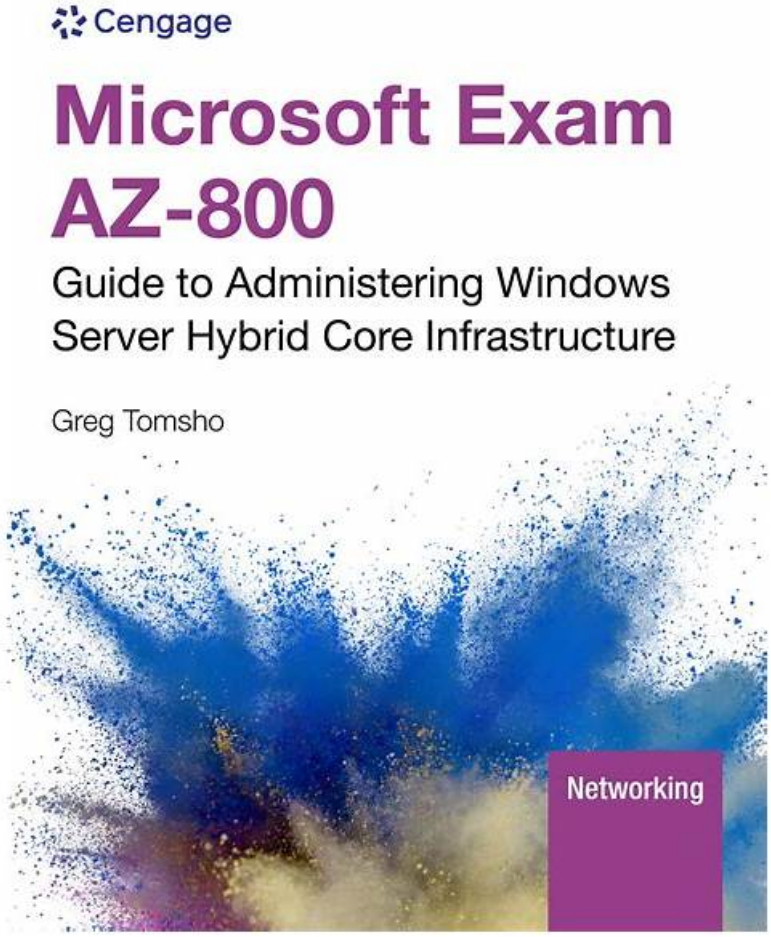

In today's society, many people are busy every day and they think about changing their status of profession. They want to improve their competitiveness in the labor market, but they are worried that it is not easy to obtain the certification of DP-800. Our study tool can meet your needs. Once you use our DP-800 exam materials, you don't have to worry about consuming too much time, because high efficiency is our great advantage. In a matter of seconds, you will receive an assessment report based on each question you have practiced on our DP-800 test material. The final result will show you the correct and wrong answers so that you can understand your learning ability so that you can arrange the learning tasks properly and focus on the targeted learning tasks with DP-800 test questions. So you can understand the wrong places and deepen the impression of them to avoid making the same mistake again.

Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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	• 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	• 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.

Right Q&A in Microsoft DP-800 Exam Questions

We are popular not only because we own the special and well-designed DP-800 exam materials but also for we can provide you with well-rounded services beyond your imagination. At the very beginning, we have an authoritative production team and our DP-800 study guide is revised by hundreds of experts, which means that you can receive a tailor-made DP-800 Study Material according to the changes in the syllabus and the latest development in theory and breakthroughs. Without doubt, our DP-800 practice torrent keep up with the latest information.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q44-Q49):

NEW QUESTION # 44
Case Study 2 - Fabrikam
Existing Environment
Azure Environment
Fabrikam has a single Azure subscription in the East US 2 Azure region. The subscription contains an Azure SQL database named DB1. DB1 contains the following tables:
* Patients
* Employees
* Procedures
* Transactions
* UsefulPrompts
* ProcedureDocuments
You store a column master key as a secret in Azure Key Vault.
You have an on-premises application named TransactionProcessing that uses a hard-coded username and password in a connection string to access DB1.
Problem Statements
Users report that after executing a long-running stored procedure named sp_UpdateProcedureForPatient, updates to the underlying data are sometimes inconsistent.
Requirements
Planned Changes
Fabrikam plans to manage all changes to Azure SQL Database objects by using source control in GitHub. Every pull request submitted to production will be validated before it can be merged.
Deployments must use the Release configuration.
Security Requirements
Fabrikam identifies the following security requirements:
* The TransactionProcessing application must use a passwordless connection to DB1.
* The Employees table contains two columns named TaxID and Salary that must be encrypted at rest.
* Auditors must have a tamper-evident history of transactions with cryptographic proof of changes to the employee data.
Database Performance Requirements
Records accessed by using sp_UpdateProcedureForPatient must NOT be changed by other transactions while the stored procedure runs.
AI Search, Embeddings, and Vector Indexing
Fabrikam identifies the following AI-related requirements:
* Queries to the ProcedureDocuments table must use Reciprocal Rank Fusion (RRF).
* Users must be able to query the data in DB1 by using prompts in Copilot in Microsoft Fabric.
* The UsefulPrompts table will store prompts that doctors can use to help diagnose patient illness by connecting to an Azure OpenAI endpoint.
Development Requirements
Fabrikam identifies the following development requirements:
* Provide the functionality to retrieve all the transactions of a given patient between two dates, showing a running total.

* Expose a Data API builder (DAB) configuration file to enable Azure services to perform the following operations over a REST API:
- Read data from the procedures table without authentication.
- Read and insert data into the Transactions table once authenticated.
- Execute the sp_UpdateProcedurePatient stored procedure.
* Provide the functionality to retrieve a list of the names of patients who underwent medical procedures during the last 30 days.
* Information for each medical procedure will be stored in a table. The table will be used with a large language model (LLM) for user querying and will have the following structure.

```
CREATE TABLE dbo.ProcedureDocuments
  (
    DocumentId INT IDENTITY PRIMARY KEY,
    SourceId NVARCHAR(200) NULL,
    Content NVARCHAR(MAX) NOT NULL,
    Embedding VECTOR(1536) NOT NULL,
    CreatedAt DATETIME2 NOT NULL DEFAULT SYSUTCDATETIME()
  );
```


DAB
You create a DAB configuration file that meets the development requirements for DB1 and includes the following entities.

You implement ProcedureDocuments to support the planned changes.
When users consume data through the Retrieval Augmented Generation (RAG) pattern, they experience data retrieval delays.
You need to improve the data retrieval performance and reduce the number of tokens per retrieval.
What should you implement?

- A. JSON content
- B. chunking
- C. a small language model (SLM)
- D. embeddings

Answer: D

Explanation:
Scenario: Fabrikam identifies the following AI-related requirements: Queries to the ProcedureDocuments table must use Reciprocal Rank Fusion (RRF).
To remedy data retrieval delays in a Retrieval Augmented Generation (RAG) pattern using Reciprocal Rank Fusion (RRF) on an Azure SQL Database table, you should use embeddings.
In a RAG architecture, retrieval delays often stem from inefficient or computationally heavy search processes. While RRF is excellent for merging results from multiple sources (like combining keyword and vector searches), the core of the speed problem typically lies in how the initial data is indexed and retrieved.
Role of Embeddings
Vector Search Acceleration: Embeddings convert text into high-dimensional vectors. Azure SQL Database can perform similarity searches on these vectors much faster than complex semantic text matching.
Hybrid Search Synergy: RRF is most effective when it fuses results from a keyword search (fast) and a vector search (powered by

embeddings). Using embeddings ensures that the "semantic" side of the retrieval is streamlined.
Pre-computation: Since embeddings are generated once during ingestion, the retrieval phase only requires a distance calculation (e.g., Cosine Similarity), which is significantly faster than real- time natural language parsing during each query.
Reference:
https://pratikbarjatya.medium.com/unlocking-the-power-of-language-with-retrieval-augmented- generation-rag-14123cc275e6

NEW QUESTION # 45
You have a SQL database in Microsoft Fabric that contains a column named Payload. Payload stores customer data in JSON documents that have the following format.

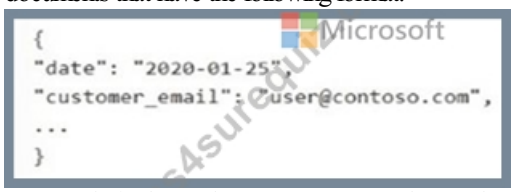

Data analysis shows that some customers have subaddressing in their email address, for example, user1+promo@contoso.com. You need to return a normalized email value that removes the subaddressing, for example, user1 +promo@contoso.com must be normalized to user1@contoso.com.
Which Transact-SQL expression should you use?

- A. REGEXP_REPLACE(JSON_VALUE(Payload, '$.customer_email'), '\+.*@', '@')
- B. REGEXP_SUBSTR(JSON_VALUE(Payload, '$.customer_email'), '