Microsoft DP-800 Practice Exam Online, DP-800 Latest Learning Material

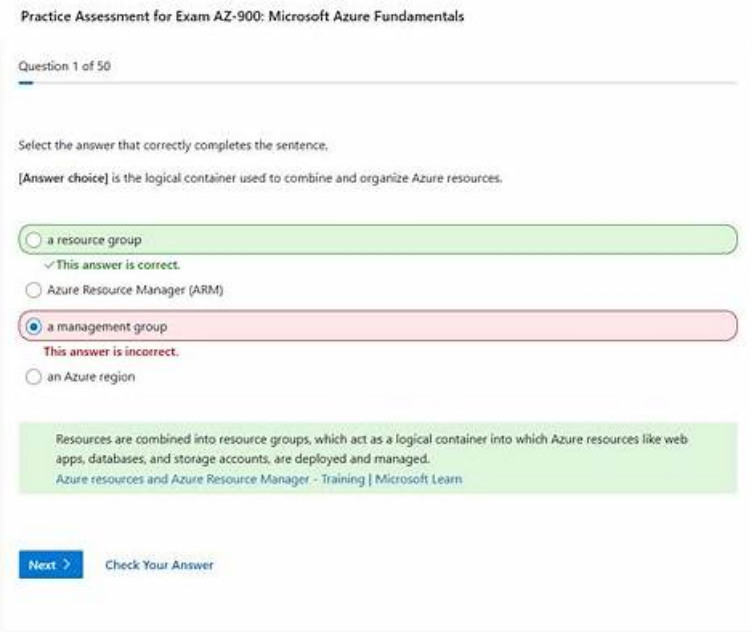

Prep4sures provides you with tri-format prep material compiled under the supervision of 90,000 Microsoft professionals from around the world that includes everything you need to pass the Microsoft DP-800 Exam on your first try. The preparation material consists of a PDF, practice test software for Windows, and a web-based practice exam. All of these preparation formats are necessary for complete and flawless preparation.

Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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 2	• 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.
Topic 3	• 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.

>> Microsoft DP-800 Practice Exam Online <<

DP-800 Latest Learning Material, DP-800 Top Questions

In today's society, there are increasingly thousands of people put a priority to acquire certificates to enhance their abilities. With a total new perspective, our DP-800 study materials have been designed to serve most of the office workers who aim at getting the DP-800 exam certification. Moreover, our DP-800 Exam Questions have been expanded capabilities through partnership with a network of reliable local companies in distribution, software and product referencing for a better development. We are helping you pass the DP-800 exam successfully has been given priority to our agenda.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q42-Q47):

NEW QUESTION # 42
You have an SDK-style SQL database project stored in a Git repository. The project targets an Azure SQL database.
The CI build fails with unresolved reference errors when the project references system objects.
You need to update the SQL database project to ensure that dotnet build validates successfully by including the correct system objects in the database model for Azure SQL Database.
Solution: Add the Microsoft.SqlServer.Dacpacs.Azure.Master NuGet package to the project.
Does this meet the goal?

- A. No
- B. Yes

Answer: B

Explanation:
This does meet the goal. Microsoft documents that SDK-style SQL projects can add the master.dacpac database reference as a package reference , and for Azure SQL Database the correct package is the Azure- specific master DACPAC package. The Azure SQL system DACPACs are available through NuGet, and this is the recommended way to include the right system objects in the database model for dotnet build validation.
So for an SDK-style SQL database project that targets Azure SQL Database, adding Microsoft.SqlServer.
Dacpacs.Azure.Master is the correct fix for unresolved references to system objects.

NEW QUESTION # 43
You have an Azure SQL database named SalesDB on a logical server named sales-sql01.
You have an Azure App Service web app named OrderApi that connects to SalesDB by using SQL authentication.
You enable a user-assigned managed identity named OrderApi-Id for OrderApi.
You need to configure OrderApi to connect to SalesDB by using Microsoft Entra authentication. The managed identity must have read and write permissions to SalesDB.
Which Transact-SQL statements should you run in SalesDB?

- A. CREATE LOGIN [OrderApi-Id] WITH PASSWORD = ' P@ssw0rd! ' ;
  ALTER SERVER ROLE sysadmin ADD MEMBER [OrderApi-Id];
- B. CREATE USER [OrderApi-Id] WITH PASSWORD = ' P@ssw0rd! ' ;
  ALTER ROLE db_datareader ADD MEMBER [OrderApi-Id];
  ALTER ROLE db_datawriter ADD MEMBER [OrderApi-Id];
- C. CREATE LOGIN [OrderApi-Id] FROM EXTERNAL PROVIDER;
  ALTER ROLE db_datareader ADD MEMBER [OrderApi-Id];
  ALTER ROLE db_datawriter ADD MEMBER [OrderApi-Id];
- D. CREATE USER [OrderApi-Id] FROM EXTERNAL PROVIDER;
  ALTER ROLE db_datareader ADD MEMBER [OrderApi-Id];
  ALTER ROLE db_datawriter ADD MEMBER [OrderApi-Id];

Answer: D

Explanation:
For an Azure App Service using a user-assigned managed identity to connect to Azure SQL Database with Microsoft Entra authentication , the required database-side step is to create a database user from the external provider , then grant the needed database roles. Microsoft's Azure SQL documentation for managed identities states that to let a managed identity access the target database, you create a SQL user for that identity by using:
CREATE USER [< identity-name >] FROM EXTERNAL PROVIDER;
and then assign the appropriate roles.
That makes db_datareader and db_datawriter the right role grants here, because the requirement says the identity must have read and write permissions to SalesDB.
The other options are incorrect:
* A uses CREATE LOGIN ... FROM EXTERNAL PROVIDER, which is not the right choice for this Azure SQL Database scenario; the documented pattern is to create a database user from the external provider.
* B and D create SQL-authentication principals with passwords, which does not meet the Microsoft Entra managed-identity requirement.

* D also grants sysadmin, which is a server-level overgrant and not appropriate for the stated read/write requirement.

NEW QUESTION # 44
Vou have a Microsoft Fabric workspace named Workspace1 that contains a SQL database named SalesDB and an API for GraphQL tern named SalesApi.
You have a Microsoft Entra group named SqlUsers.
From Workspace1, you assign permission to SalesApi as shown in the following exhibit.
The connection to SalesDB has the connectivity option configured as shown in the following exhibit.
SqlUsers has the Viewer role for Workspace1.
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 members of SqlUsers can modify the data in SalesDB via SalesApi. # No
* The members of SqlUsers can view the data in SalesDB via SalesApi. # No
* The members of SqlUsers can change the field mappings of SalesApi. # Yes For both viewing and modifying data through SalesApi, the key missing permission is Run Queries and Mutations . Microsoft's Fabric GraphQL documentation states that callers need Execute permissions for the GraphQL API , which correspond to the Run Queries and Mutations option, and with SSO connectivity they also need appropriate permissions on the underlying data source. In the exhibit, SqlUsers has View and Edit GraphQL item selected, but Run Queries and Mutations is not selected, so members cannot query or mutate data through the API. The third statement is Yes because the group was explicitly granted View and Edit GraphQL item . That permission is the one that allows users to open and modify the GraphQL item itself, including schema-related configuration such as field mappings in the API item. The workspace Viewer role does not by itself grant query execution through the API, but the direct GraphQL item permission shown does allow editing the item.

NEW QUESTION # 45
You need to recommend a solution for the development team to retrieve the live metadata. The solution must meet the development requirements.
What should you include in the recommendation?

- A. Include the database project in the code repository.
- B. Export the database schema as a .dacpac file and load the schema into a GitHub Copilot context window.
- C. Use an MCP server
- D. Add the schema to a GitHub Copilot instruction file.

Answer: C

Explanation:
The best recommendation is to use an MCP server . In the official DP-800 study guide , Microsoft explicitly lists skills such as configuring Model Context Protocol (MCP) tool options in a GitHub Copilot session and connecting to MCP server endpoints, including Microsoft SQL Server and Fabric Lakehouse . That makes MCP the exam-aligned mechanism for enabling AI-assisted tools to work with live database context rather than static snapshots.
This also matches the stated development requirement: the team will use Visual Studio Code and GitHub Copilot and needs to retrieve live metadata from the databases . Microsoft's documentation for GitHub Copilot with the MSSQL extension explains that Copilot works with an active database connection , provides schema-aware suggestions , supports chatting with a connected database, and adapts responses based on the current database context . Microsoft also documents MCP as the standard way for AI tools to connect to external systems and data sources through discoverable tools and endpoints.
The other options do not satisfy the "live metadata" requirement as well:
* A .dacpac is a point-in-time schema artifact, not live metadata.
* A Copilot instruction file provides guidance, not live database discovery.
* Including the database project in the repository helps source control and deployment, but it still does not provide live database metadata by itself.

NEW QUESTION # 46
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: C

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;
* Al 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 VehicleTd, Lastupdatedutc, EngineStatus, BatteryHealth FROM dbo.VehicleHealthSumary where fleetld - gFleetld ORDER BV 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 Al-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 LastModif reduce column to the time of the change.
* Avoid recursive updates.
AI Search, Embedding's, and Vector indexing
The development learn 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 # 47
.....

After you enter the examination room and get the exam paper, you must be sighed that the gold content of our DP-800 learning guide is too high. Our DP-800 study materials are really magic weapon for you to quickly pass the exam. Just come and buy our DP-800 Exam Questions, then you can pass the exam by 100% success guarantee after you prapare with them for 20 to 30 hours. This data is created by our loyal customers who had bought our DP-800 training engine and passed the exam.