# Detailed Microsoft DP-800 Study Dumps, Exam DP-800 Cram Review

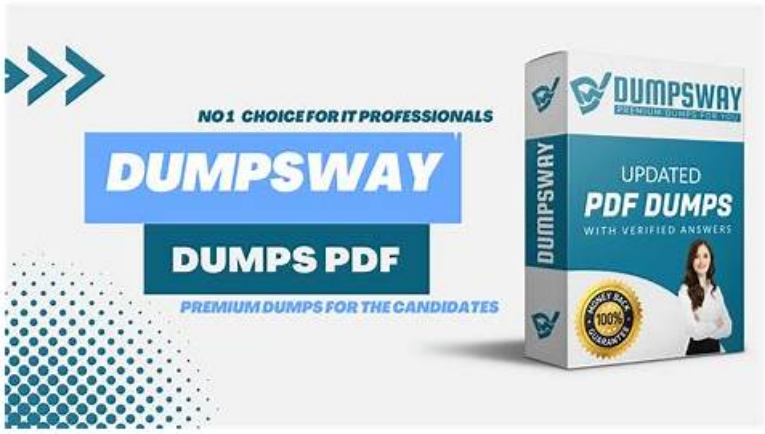


By these three versions of DP-800 practice materials we have many repeat orders in a long run. The PDF version helps you read content easier at your process of studying with clear arrangement, and the PC Test Engine version of DP-800 practice materials allows you to take stimulation exam to check your process of exam preparing, which support windows system only. Moreover, there is the APP version of DP-800 practice materials, you can learn anywhere at any time with it at your cellphones without the limits of installation.

## 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 | • 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 | • 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<br>• CD pipelines using SQL Database Projects, and integrating SQL solutions with Azure services including Data API builder and monitoring tools. |

**>> Detailed Microsoft DP-800 Study Dumps <<**

## Microsoft DP-800 Exam Questions - The Advantages of VCE4Dumps Preparation Material

Some of our customers are white-collar workers with no time to waste, and need a Microsoft certification urgently to get their promotions, meanwhile the other customers might aim at improving their skills. So we try to meet different requirements by setting different versions of our DP-800 question and answers. The special one is online DP-800 engine version. As an online tool, it is convenient and easy to study, supports all Web Browsers and system including Windows, Mac, Android, iOS and so on. You can apply this version of DP-800 exam questions on all eletric devices.

## Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q36-Q41):

**NEW QUESTION # 36**
You have an Azure SQL database that supports the OLTP workload of an order-processing application.
During a 10-minute incident window, you run a dynamic management view query and discover the following:
Session 72 is sleeping with open_transaction_count = 1.
Multiple other sessions show blocking_session_id = 72 in sys.dm_exec_requests.
sys.dm_exec_input_buffer(72, NULL) returns only BEGIN TRANSACTION UPDATE Sales.Orders.
Users report that updates to Sales.Orders intermittently time out during the incident window. The timeouts stop only after you manually terminate session 72.
What is a possible cause of the blocking?

- A. An explicit transaction was started but not committed or rolled back.
- B. A long-running SELECT statement is blocking writers.
- C. A lock escalation occurred.
- D. Session 72 caused a deadlock.

**Answer: A**

Explanation:
The best explanation is an open explicit transaction . During the incident, session 72 was sleeping but still had open_transaction_count = 1 , and sys.dm_exec_input_buffer(72, NULL) showed only BEGIN TRANSACTION UPDATE Sales.Orders. That pattern indicates the session executed an update inside an explicit transaction and then remained idle without committing or rolling back , while still holding locks.
Other sessions showing blocking_session_id = 72 is the expected symptom of that situation. Microsoft explains that blocking occurs when one session holds a lock on a resource and another session requests a conflicting lock, and sleeping sessions can continue to block if they retain locks through an open transaction.
This also fits the observed behavior that the timeouts stopped only after session 72 was terminated . Killing the session would roll back the active transaction and release the locks, allowing waiting updates to continue.
That is much more consistent with an uncommitted transaction than with a deadlock, because deadlocks are normally detected and one session is chosen as the victim automatically rather than persisting until manual termination.

**NEW QUESTION # 37**
Vou have a SQL database in Microsoft Fabric that contains a nvarchar(max) column named MessageText. An ID is always contained within the first paragraph of MessageText.
You need to write a Transact SQL query that uses REGEXP_SUBSTR to extract the ID from MessageText.
What should you include in the query?

- A. Add a collate Lati1_General_CS_AS clause to MessageText before calling regexp_SUBSTR.
- B. Apply STRING_ESCAPE(MessageText, " json' > before calling REGEXP_SUBSTR.
- C. Cast Messagelext to nvarchar(4000) before calling REGEXP SUBSIR.
- D. Run TRV_CONVERT(varchar(Bax), MessageText) before calling REGEXP_SUBSTR.

**Answer: C**

Explanation:
Microsoft documents REGEXP_SUBSTR for Transact-SQL with the string_expression parameter as supporting character string types char, nchar, varchar, and nvarchar. For the regex functions, support for LOB types such as varchar(max) and nvarchar(max) is specifically called out for REGEXP_LIKE , REGEXP_COUNT , and REGEXP_INSTR up to 2 MB, but that support note is not listed for REGEXP_SUBSTR in the surfaced documentation. In exam terms, the safe and expected approach is to cast the nvarchar(max) column to nvarchar(4000) before calling REGEXP_SUBSTR.
This also fits the scenario detail that the ID is always contained within the first paragraph of MessageText.
Since the needed value is near the start of the text, narrowing the input to a non-LOB string type such as nvarchar(4000) is sufficient and avoids incompatibility concerns with nvarchar(max).
The other options are not appropriate:
* A STRING_ESCAPE(..., ' json ' ) is for JSON escaping, not regex extraction.
* C adding a case-sensitive collation changes comparison behavior, but it is not the required fix for REGEXP_SUBSTR on nvarchar(max).
* D TRY_CONVERT(varchar(max), ...) still leaves a MAX type and also risks unnecessary Unicode loss.

**NEW QUESTION # 38**

You have an Azure SQL database that contains a table named stores, stores contains a column named description and a vector column named embedding.
You need to implement a hybrid search query that meets the following requirements:
* Uses full-text search on description for the keyword portion
* Returns the top 20 results based on a combined score that uses a weighted formula of 60% vector distance and 40% full-text rank
How should you configure the query components? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Answer:**

Explanation:

Explanation:

For the vector portion, the correct choice is VECTOR_DISTANCE and order by distance ascending . The requirement is to build a combined weighted formula using the actual vector distance. Microsoft documents that VECTOR_DISTANCE returns the exact distance between two vectors. Since lower distance means greater similarity, ascending distance is the right direction for ranking. VECTOR_SEARCH is for ANN retrieval, but this hotspot specifically asks for a weighted formula based on distance , so VECTOR_DISTANCE is the appropriate operator.
For the keyword portion, the correct choice is CONTAINSTABLE on description and return ranked matches . Microsoft documents that CONTAINSTABLE returns a RANK column from 0 through 1000 , which is exactly what is needed for weighted scoring in a hybrid formula.
For the final ranking expression, the best choice is order by (distance * 0.6) + ((1.0 - RANK/1000.0) * 0.4) .
This works because vector distance is a lower-is-better metric, while full-text RANK is a higher-is-better metric. Dividing RANK by 1000 normalizes it to the documented range, and subtracting from 1.0 converts it into a lower-is-better term so both components can be combined consistently in one ascending score. This final step is a sound inference based on Microsoft's documented distance semantics and full-text rank range.

**NEW QUESTION # 39**
You need to meet the development requirements for the FeedbackJson column How should you complete the Transact SQL query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

**Answer:**

Explanation:

Explanation:

JSON_VALUE(f.FeedbackJson, ' $.text ' ) AS FeedbackText
CONTAINS(FeedbackJson, @Keyword)
SimilarityScore
These three selections are the correct way to complete the query because they align exactly with the stated requirements for the FeedbackJson column.
First, to extract the customer feedback text from the JSON document , the correct expression is JSON_VALUE(f.FeedbackJson, ' $.text ' ) AS FeedbackText . Microsoft documents that JSON_VALUE is used to extract a scalar value from JSON, while JSON_QUERY is used for returning an object or array .
Since $.text is the textual feedback string, JSON_VALUE is the correct function.
Second, to filter rows where the JSON text contains a keyword , the best choice is CONTAINS (FeedbackJson, @Keyword) . The scenario explicitly states that FeedbackJson already has a full-text index
, and Microsoft documents that CONTAINS is the full-text predicate used in the WHERE clause to search full-text indexed character data. That makes it more appropriate than using EDIT_DISTANCE for keyword filtering.
Third, to order the results by similarity score, highest first , the correct item is SimilarityScore in the ORDER BY clause, which would be paired with DESC in the query. This matches the requirement to sort by the computed fuzzy similarity value. The DP-800 study guide specifically includes writing queries that use fuzzy string matching functions such as EDIT_DISTANCE, which supports the earlier computed SimilarityScore expression in the query.

**NEW QUESTION # 40**
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 # 41**
.....

Especially for those students who are headaches when reading a book, DP-800 study tool is their gospel. Because doing exercises will make it easier for one person to concentrate, and at the same time, in the process of conducting a mock examination to test yourself, seeing the improvement of yourself will makes you feel very fulfilled and have a stronger interest in learning. DP-800 Guide Torrent makes your learning process not boring at all.