



**Microsoft**

**DP-800 Exam**

Developing AI-Enabled Database Solutions

**Thank you for Downloading DP-800 exam PDF Demo**

You can also try our DP-800 practice exam software

**Download Free Demo**

<https://www.certprep.io/DP-800.html>

QUESTION & ANSWERS  
**DEMO VERSION**  
(LIMITED CONTENT)

# Version: 4.0

---

**Question: 1**

---

You need to generate embeddings to resolve the issues identified by the analysts. Which column should you use?

- A. vehicleLocation
- B. incidentDescription
- C. incidentType
- D. SeverityScore

---

**Answer: B**

---

Explanation:

The correct column to use for generating embeddings is incidentDescription because embeddings are intended to represent the semantic meaning of rich textual content, not simple categorical, numeric, or location-only values. Microsoft's DP-800 study guide explicitly includes skills such as identifying which columns to include in embeddings, generating embeddings, and implementing semantic vector search for scenarios where users need to find similar records based on meaning rather than exact matches.

In this scenario, analysts report that it is difficult to find similar incidents based on details such as weather, traffic conditions, and location. Those are descriptive context elements that are typically captured in a free-text incident description field. An embedding generated from incidentDescription can encode the semantic relationships among these narrative details, making it suitable for similarity search, semantic search, and RAG retrieval. Microsoft documentation on vectors and embeddings explains that embeddings are generated from text data and then stored for vector search to find semantically related items.

The other options are weaker choices:

vehicleLocation is too narrow and usually better handled with geospatial filtering, not embeddings.

incidentType is likely categorical and too low in semantic richness.

SeverityScore is numeric and not appropriate as the primary source for semantic embeddings.

Microsoft also notes that when multiple useful attributes exist, you can either embed each text column separately or concatenate relevant text fields into one textual representation before

generating the embedding. But among the options given, the best and most exam-aligned answer is the textual narrative column: `incidentDescription`.

---

**Question: 2**

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

---

**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.

---

**Question: 3**

---

You need to recommend a solution that will resolve the ingestion pipeline failure issues. Which two actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Enable snapshot isolation on the database.
- B. Use a trigger to automatically rewrite malformed JSON.
- C. Add foreign key constraints on the table.
- D. Create a unique index on a hash of the payload.
- E. Add a check constraint that validates the JSON structure.

---

**Answer: D, E**

---

Explanation:

The two correct actions are D and E because the ingestion failures are caused by malformed JSON and duplicate payloads, and these two controls address those two problems directly. Microsoft's JSON documentation states that SQL Server and Azure SQL support validating JSON with ISJSON, and Microsoft specifically recommends using a CHECK constraint to ensure JSON text stored in a column is properly formatted.

For the duplicate-payload issue, creating a unique index on a hash of the payload is the appropriate design. Microsoft documents using hashing functions such as HASHBYTES to hash column values, and SQL Server allows a deterministic computed column to be used as a key column in a UNIQUE constraint or unique index. That makes a persisted hash-based computed column plus a unique index a practical and exam-consistent way to reject duplicate payloads efficiently.

The other options do not solve the stated root causes:

Snapshot isolation addresses concurrency behavior, not malformed JSON or duplicate payload detection.

A trigger to rewrite malformed JSON is not the right integrity control and is brittle.

Foreign key constraints enforce referential integrity, not JSON validity or duplicate-payload prevention

---

**Question: 4**

---

HOTSPOT

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 Area**

```

SELECT
    f.FeedbackId,
    f.VehicleId,
    CONTAINS(FeedbackJson, @keyword)
    EDIT_DISTANCE( JSON_VALUE(f.FeedbackJson, '$.details.comment'), @Keyword) < 3
    EDIT_DISTANCE(JSON_VALUE(f.FeedbackJson, '$.text'), @Keyword) < 3
    JSON_QUERY(f.FeedbackJson, '$.text', @KnownIssueDescription) AS FeedbackText
    JSON_VALUE(f.FeedbackJson, '$.text') AS FeedbackText
    SimilarityScore
FROM
    dbo.CustomerFeedback f
WHERE
    CONTAINS(FeedbackJson, @Keyword)
    EDIT_DISTANCE( JSON_VALUE(f.FeedbackJson, '$.details.comment'), @Keyword) < 3
    EDIT_DISTANCE(JSON_VALUE(f.FeedbackJson, '$.text'), @Keyword) < 3
    JSON_QUERY(f.FeedbackJson, '$.text', @KnownIssueDescription) AS FeedbackText
    JSON_VALUE(f.FeedbackJson, '$.text') AS FeedbackText
    SimilarityScore
    
```

---

**Answer:**

---

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.

---

**Question: 5**

---

DRAG DROP

You need to meet the database performance requirements for maintenance data

How should you complete the Transact-SQL code? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

**Values**

- i.MaintenanceId IS NOT NULL
- m.LastModifiedUtc <> i.LastModifiedUtc
- m.MaintenanceId = i.MaintenanceId
- m.VehicleId = i.VehicleId

**Answer Area**

```

CREATE TRIGGER dbo.trgMaintenanceEvents_UpdateTimestamp
ON dbo.MaintenanceEvents
AFTER UPDATE
AS
BEGIN
UPDATE m
SET LastModifiedUtc = SYSUTCDATETIME()
FROM dbo.MaintenanceEvents m
INNER JOIN inserted i
ON [Value]
WHERE [Value]
END;
GO
        
```

---

**Answer:**

---

Explanation:

ON → m.maintenancelid = i.maintenancelid

WHERE → m.LastModifiedUtc <> i.LastModifiedUtc

The correct drag-and-drop completion is:

ON m.maintenancelid = i.maintenancelid

WHERE m.LastModifiedUtc <> i.LastModifiedUtc

This satisfies the requirement to ensure that when a row in MaintenanceEvents changes, the corresponding LastModifiedUtc value is updated to the current system time, while also helping avoid unnecessary repeat updates.

The inserted pseudo-table in a SQL Server AFTER UPDATE trigger contains the rows that were just updated. To update the matching row in the base table correctly, the trigger must join the target table row to the corresponding row in inserted by the table’s primary key. In this schema, MaintenanceId is the primary key for MaintenanceEvents, so the correct join is m.maintenancelid = i.maintenancelid. Joining on VehicleId would be incorrect because multiple maintenance rows could exist for the same vehicle, which could update unintended rows. Microsoft’s trigger documentation explains that inserted and deleted are used to work with the affected rows and that multi-row logic should be based on proper key matching.

The WHERE m.LastModifiedUtc <> i.LastModifiedUtc predicate is used to prevent the trigger from re-updating rows where the timestamp already matches the value in inserted. That reduces redundant writes and supports the requirement to avoid recursive or repeated update behavior. In practice, this

means the trigger updates only rows whose current stored timestamp differs from the just-updated version. This is the exam-appropriate pattern for a self-updating timestamp column in an AFTER UPDATE trigger.

## Thank You for trying DP-800 PDF Demo

To try our DP-800 practice exam software visit link below

<https://www.certprep.io/DP-800.html>

## Start Your DP-800 Preparation

Use Coupon “**SAVE50**” for extra 50% discount on the purchase of Practice Test Software. Test your DP-800 preparation with actual exam questions.