

DP-800 Demotesten & DP-800 Vorbereitungsfragen



Mit der Microsoft DP-800 Zertifizierungsprüfung werden Sie sicher bessere Berufsaussichten haben. Die Microsoft DP-800 Zertifizierungsprüfung kann nicht nur Ihre Fertigkeiten, sondern auch Ihre Zertifikate und Fachkenntnisse beweisen. Die den Schulungsunterlagen zur Microsoft DP-800 Zertifizierungsprüfung von ITZert sind eine von der Praxis bewährte Software. Mit ihr können Sie eine bessere Theorie bekommen. Vorm Kauf können Sie eine kostenlose Probeversion bekommen. So kennen Sie die Qualität unserer Prüfungsmaterialien. ITZert ist Ihnen die beste Wahl.

Die Microsoft DP-800 Prüfung macht man wirklich besorgt. Vielleicht vertragen Sie nicht mehr die große Menge von

Prüfungsunterlagen, dann lassen Sie Microsoft DP-800 Prüfungssoftware von ITZert Ihnen helfen, die Belastungen zu erleichtern! Unsere professionelle IT-Profis haben die anspruchsvolle Microsoft DP-800 Prüfungssoftware entwickelt dadurch, dass die komplizierten Test-Bank geordnet und die Schwerpunkte der Prüfungen in den letzten Jahren analysiert haben. Trotzdem aktualisieren wir die Microsoft DP-800 Prüfungsunterlagen immer weiter. Innerhalb einem Jahr nach Ihrem Kauf geben wir Ihnen sofort Bescheid, wenn die Microsoft DP-800 aktualisiert hat.

>> DP-800 Demotesten <<

DP-800 Vorbereitungsfragen, DP-800 Prüfungsvorbereitung

Wir ITZert bieten Ihnen verschiedene Unterlagenversionen, die Ihre Nutzung erleichtern können. Die PDF-Versionen können das Lesen erleichtern und Ihnen die aktuellen Microsoft DP-800 Prüfungsfragen zeigen, Die Software-Versionen sind die Simulationssoftwares, die Ihre Vorbereitungssituation auf jeden Fall testen. Wenn Sie wissen wollen, ob Sie sich für Microsoft DP-800 Prüfung gut bereit sind, können Sie helfen, Ihre Stärke und Schwäche ganz schnell finden, um Ihren nächsten Lernplan zu erstellen.

Microsoft Developing AI-Enabled Database Solutions DP-800 Prüfungsfragen mit Lösungen (Q37-Q42):

37. Frage

You have an Azure SQL database that has Query Store enabled
Query Performance Insight shows that one stored procedure has the longest runtime. The procedure runs the following parameterized query.

```

CREATE OR ALTER PROCEDURE dbo.GetRecentOrders
    @customerId int,
    @sinceDate datetime2
AS
SELECT TOP (50)
    o.OrderId, o.OrderDate, o.Status, o.TotalAmount
FROM dbo.Orders AS o
WHERE o.CustomerId = @CustomerId
    AND o.OrderDate >= @sinceDate
ORDER BY o.OrderDate DESC;
    
```

The dbo.orders table has approximately 120 million rows. Customer-id is highly selective, and orderOate is used for range filtering and sorting.

You have the following indexes:

- * Clustered index: PK_Orders on (OrderId)
 - * Nonclustered index: IX_Orders_order-Date on (OrderDate) with no included columns
- An actual execution plan captured from Query Store for slow runs shows the following:
- * An index seek on ixordersorderDate followed by a Key Lookup (Clustered) on PKOrders for customerid, status, and TotalAmount
 - * A sort operator before top (50), because the results are ordered by orderDate DESC
- For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
To avoid the explicit sort for the query, create a nonclustered index on (CustomerId, OrderDate DESC) that includes (Status, TotalAmount).	<input type="radio"/>	<input type="radio"/>
To eliminate the sort and make the query use an ordered seek, add CustomerId as an included column to IX_Orders_OrderDate.	<input type="radio"/>	<input type="radio"/>
The plan indicates a bottleneck from a suboptimal query plan, rather locking or blocking.	<input type="radio"/>	<input type="radio"/>

Antwort:

Begründung:

ANSWER AREA

Statements	Yes	No
To avoid the explicit sort for the query, create a nonclustered index on (CustomerId, OrderDate DESC) that includes (Status, TotalAmount).	<input type="checkbox"/>	<input type="checkbox"/>
To eliminate the sort and make the query use an ordered seek, add CustomerId as an included column to IX_orders_OrderDate.	<input type="checkbox"/>	<input checked="" type="checkbox"/>
The plan indicates a bottleneck from a suboptimal query plan, rather locking or blocking.	<input type="checkbox"/>	<input type="checkbox"/>

Explanation:

Answer Area

Statements	Yes	No
To avoid the explicit sort for the query, create a nonclustered index on (CustomerId, OrderDate DESC) that includes (Status, TotalAmount).	<input checked="" type="checkbox"/>	<input type="checkbox"/>
To eliminate the sort and make the query use an ordered seek, add CustomerId as an included column to IX_orders_OrderDate.	<input type="checkbox"/>	<input checked="" type="checkbox"/>
The plan indicates a bottleneck from a suboptimal query plan, rather locking or blocking.	<input checked="" type="checkbox"/>	<input type="checkbox"/>

The first statement is Yes . The query filters on CustomerId, applies a range predicate on OrderDate, and sorts by OrderDate DESC. Microsoft's index design guidance recommends putting equality predicates first in the key, followed by columns used for ordering/range access, because the order of key columns determines seek and sort support. A nonclustered index on (CustomerId, OrderDate DESC) can support an ordered seek for this query and avoid the explicit sort. Including Status and TotalAmount helps cover the query, and OrderId is already available because the clustered key is stored with nonclustered index rows.

The second statement is No . Adding CustomerId as an included column to IX_Orders_OrderDate does not make it part of the index's navigational structure. Microsoft states that included columns are nonkey columns used to cover queries; they do not provide the seek and ordering characteristics that key columns do. So an index keyed only on OrderDate still is not the right ordered access path for WHERE CustomerId =

@CustomerId ... ORDER BY OrderDate DESC.

The third statement is Yes . The described actual plan shows an index seek on the wrong access path for the workload, followed by clustered key lookups and an explicit sort before TOP (50). That is characteristic of a suboptimal query/index plan . Query Store and Query Performance Insight are designed to surface plan- related performance regressions, while locking/blocking problems are typically identified through waits

/DMVs and blocking-session indicators, not from a plan shape like seek + lookup + sort alone.

38. Frage

You have an Azure SQL database named ProductsDB.

You deploy Data API builder (DAB) to Azure Container Apps by using the mcr.microsoft.com/azure-databases/data-api-builder:latest image.

The container app has the following configurations:

- Secrets: mssql-connection-string, dab-config-base64
- Environment variables:
 - MSSQL_CONNECTION_STRING=secretref:mssql-connection-string
 - DAB_CONFIG_BASE64=secretref:dab-config-base64
- Ingress: External on port 5000

Users report that the /health endpoint returns a healthy response, but all requests that query an entity named Products fail and generate a connection error.

You confirm that the SQL login in the connection string is correct and the database exists.

You need to ensure that the container app can establish connections to the Azure SQL logical server without changing the container app deployment settings or the DAB configuration file.

What should you do on the Azure SQL logical server?

- A. Create a firewall rule that allows a start and end IP address of 0.0.0.0.
- B. Run DBCC CHECKDB on ProductsDB.
- C. Enable FORCE_LAST_GOOD_PLAN automatic tuning for ProductsDB.
- D. Create an auto-failover group for ProductsDB.

Antwort: A

Begründung:

Even if the login credentials are correct, Azure SQL Database blocks all incoming traffic by default. Since your Container App is

returning a healthy response for the /health endpoint (which is internal to the DAB engine) but failing on entity queries (which require a database hit), the network handshake is being rejected at the SQL Server firewall level.

To fix this connection error without changing the container app or the DAB configuration file, you must enable the Azure SQL Server firewall to allow Azure services.

Setting the start and end IP address to 0.0.0.0 in an Azure SQL Database firewall rule is a specific configuration that enables the "Allow Azure services and resources to access this server" setting.

Primary Use Case

In your scenario, this rule allows your Azure Container Apps (ACA) to communicate with your Azure SQL Database over the Azure backbone network.

Connectivity: It permits any traffic originating from within the Azure boundary to reach the database.

Simplicity: You do not need to track or white-list the specific outbound IP addresses of your Container App, which can change if the app scales or restarts.

Internal Routing: Traffic stays within the Azure network rather than routing out to the public internet and back in.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/database/firewall-configure>

<https://stackoverflow.com/questions/54599813/how-to-enable-the-access-to-azure-services-in-my-azure-sql-database-server>

39. Frage

Case Study 1 - Contoso

Existing Environment

Azure 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.

Table names	Column names
CustomerFeedback	<ul style="list-style-type: none"> FeedbackId (int) (primarykey) FeedbackJson (nvarchar (max))
Fleets	<ul style="list-style-type: none"> FleetId (int) (primarykey) FleetName (nvarchar(100)) Description (nvarchar(256))
MaintenanceEvents	<ul style="list-style-type: none"> MaintenanceId (int) (primarykey) VehicleId (int) LastModifiedUTC (datetime2) Description (nvarchar(256))
SupportTickets	<ul style="list-style-type: none"> TicketId (int) (primarykey) FleetId (int) CreatedUtc (datetime2)
UserAccounts	<ul style="list-style-type: none"> UserId (int) (primarykey) UserPrincipalName (nvarchar(256)) JobRole (nvarchar(256)) StartDate (datetime2)
VehicleIncidentReports	<ul style="list-style-type: none"> IncidentId (int) (primarykey) VehicleId (int) FleetId (int) IncidentType (nvarchar(50)) VehicleLocation (nvarchar(200)) IncidentDescription (nvarchar(max)) SeverityScore (int)
Vehicles	<ul style="list-style-type: none"> VehicleId (int) (primarykey) VIN ((nvarchar(50)) VehicleDescription (nvarchar(256))
VehicleHealthSummary	<ul style="list-style-type: none"> VehicleId (int) (primarykey) FleetId (int) Summary (nvarchar(2000)) LastUpdatedUtc (datetime2) EngineStatus [bit] EngineStatusLastUpdatedUtc (datetime2) BatteryHealth (int) Embeddings (vector (1536))

The FeedbackJson column has a full-text index and stores JSON documents in the following format.

```
{
  "text": "The battery drains too fast when driving uphill.",
  "category": "Battery",
  "metadata": {
    "appVersion": "5.2.1",
    "device": "Android",
    "language": "en-GB"
  }
}
```



The support staff at Contoso never has the UNMASK permission.

Problem Statements

Contoso is deploying a new Azure SQL database that will become the authoritative data store for the following:

* AI 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 VehicleId, LastUpdatedUtc, EngineStatus, BatteryHealth
FROM dbo.VehicleHealthSummary
WHERE FleetId = @FleetId
ORDER BY LastUpdatedUtc DESC;
```



You review the execution plan and discover that the plan shows a clustered index scan.

VehicleIncidentReportsoften contains details about the weather, traffic conditions, and location. Analysts report that it is difficult to find similar incidents based on these details.

Requirements

Planned Changes

Contoso wants to modernize Fleet Intelligence Platform to support AI-powered semantic search over incident reports.

Security Requirements

Contoso identifies the following security requirements:

- * Restrict the support staff from viewing Personally Identifiable Information (PII) data, which is full email addresses and phone numbers.
- * Enforce row-level filtering so that analysts see only incidents for the fleets to which they are assigned. The analysts can be assigned to multiple fleets.

Database Performance and 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 accuracyJSON 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 MaintenanceEventstable updates the corresponding value in the LastModifiedUtccolumn to the time of the change.
- * Avoid recursive updates.

AI Search, Embeddings, and Vector Indexing

Contoso plans to implement semantic search over incident data to meet the following requirements:

- * Embeddings must be stored in dedicated Azure SQL Database tables.
- * Embeddings must be generated from rich natural language fields.
- * Chunking must preserve semantic coherence.
- * Hybrid search must combine the following:

- Vector similarity
- Keyword filtering or boosting

Development Requirements

The development team 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 FeedbackJsoncolumn of the CustomerFeedbacktable:

- * 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.

Hotspot Question

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

EDIT_DISTANCE_SIMILARITY(
Microsoft

JSON_VALUE(f.FeedbackJson, '\$.text'),

@KnownIssueDescription

) AS 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

ORDER BY

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

DESC;

Antwort:

Begründung:

```

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

```

    EDIT_DISTANCE_SIMILARITY(
        JSON_VALUE(f.FeedbackJson, '$.text'),
        @KnownIssueDescription
    ) AS 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

ORDER BY

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

40. Frage

Case Study 1 - Contoso

Existing Environment

Azure 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.

Table names	Column names
CustomerFeedback	<ul style="list-style-type: none"> FeedbackId (int) (primarykey) FeedbackJson (nvarchar (max))
Fleets	<ul style="list-style-type: none"> FleetId (int) (primarykey) FleetName (nvarchar(100)) Description (nvarchar(256))
MaintenanceEvents	<ul style="list-style-type: none"> MaintenanceId (int) (primarykey) VehicleId (int) LastModifiedUTC (datetime2) Description (nvarchar(256))
SupportTickets	<ul style="list-style-type: none"> TicketId (int) (primarykey) FleetId (int) CreatedUtc (datetime2)
UserAccounts	<ul style="list-style-type: none"> UserId (int) (primarykey) UserPrincipalName (nvarchar(256)) JobRole (nvarchar(256)) StartDate (datetime2)
VehicleIncidentReports	<ul style="list-style-type: none"> IncidentId (int) (primarykey) VehicleId (int) FleetId (int) IncidentType (nvarchar(50)) VehicleLocation (nvarchar(200)) IncidentDescription (nvarchar(max)) SeverityScore (int)
Vehicles	<ul style="list-style-type: none"> VehicleId (int) (primarykey) VIN ((nvarchar(50)) VehicleDescription (nvarchar(256))
VehicleHealthSummary	<ul style="list-style-type: none"> VehicleId (int) (primarykey) FleetId (int) Summary (nvarchar(2000)) LastUpdatedUtc (datetime2) EngineStatus [bit] EngineStatusLastUpdatedUtc (datetime2) BatteryHealth (int) Embeddings (vector (1536))

The FeedbackJson column has a full-text index and stores JSON documents in the following format.

```
{
  "text": "The battery drains too fast when driving uphill.",
  "category": "Battery",
  "metadata": {
    "appVersion": "5.2.1",
    "device": "Android",
    "language": "en-GB"
  }
}
```



The support staff at Contoso never has the UNMASK permission.

Problem Statements

Contoso is deploying a new Azure SQL database that will become the authoritative data store for the following:

* AI 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 VehicleId, LastUpdatedUtc, EngineStatus, BatteryHealth
FROM dbo.VehicleHealthSummary
WHERE FleetId = @FleetId
ORDER BY LastUpdatedUtc DESC;
```



You review the execution plan and discover that the plan shows a clustered index scan.

VehicleIncidentReportsoften contains details about the weather, traffic conditions, and location. Analysts report that it is difficult to find similar incidents based on these details.

Requirements

Planned Changes

Contoso wants to modernize Fleet Intelligence Platform to support AI-powered semantic search over incident reports.

Security Requirements

Contoso identifies the following security requirements:

- * Restrict the support staff from viewing Personally Identifiable Information (PII) data, which is full email addresses and phone numbers.
- * Enforce row-level filtering so that analysts see only incidents for the fleets to which they are assigned. The analysts can be assigned to multiple fleets.

Database Performance and 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 accuracyJSON 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 MaintenanceEventstable updates the corresponding value in the LastModifiedUtccolumn to the time of the change.
- * Avoid recursive updates.

AI Search, Embeddings, and Vector Indexing

Contoso plans to implement semantic search over incident data to meet the following requirements:

- * Embeddings must be stored in dedicated Azure SQL Database tables.
- * Embeddings must be generated from rich natural language fields.
- * Chunking must preserve semantic coherence.
- * Hybrid search must combine the following:

- Vector similarity
- Keyword filtering or boosting

Development Requirements

The development team 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 FeedbackJsoncolumn of the CustomerFeedbacktable:

- * 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.

Drag and Drop Question

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	Answer Area
<input type="text" value="i.MaintenanceId IS NOT NULL"/>	<pre> 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 <input type="text"/> WHERE <input type="text"/> END; GO </pre>
<input type="text" value="m.LastModifiedUtc <> i.LastModifiedUtc"/>	
<input type="text" value="m.MaintenanceId = i.MaintenanceId"/>	
<input type="text" value="m.VehicleId = i.VehicleId"/>	

Antwort:

Begründung:

Values	Answer Area
<input type="text" value="i.MaintenanceId IS NOT NULL"/>	<pre> 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 <input type="text" value="m.VehicleId = i.VehicleId"/> WHERE <input type="text" value="m.LastModifiedUtc <> i.LastModifiedUtc"/> END; GO </pre>
<input type="text" value="m.LastModifiedUtc <> i.LastModifiedUtc"/>	
<input type="text" value="m.MaintenanceId = i.MaintenanceId"/>	
<input type="text" value="m.VehicleId = i.VehicleId"/>	

41. Frage

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. Regenerate embeddings for the existing rows.
- B. Create a vector index on `dbo.Products.Embedding`.

- C. Convert the Embedding column to nvarchar(max).
- D. Normalize the vector lengths before storing new embeddings.

Antwort: B

Begründung:

To ensure your applications can transition models without runtime errors while using VECTOR_SEARCH, you must first define a Vector Index that explicitly identifies the dimensions and distance metric.

Since you are moving from text-embedding-ada-002 to text-embedding-3-small, both models default to 1536 dimensions, which matches your existing column definition. To create the index as the first step, use the following SQL:

```
CREATE VECTOR INDEX idx_embedding ON YourTableName (Embedding)
WITH ( DISTANCE_METRIC = 'COSINE' );
```

Use code with caution.

Why this works:

Schema Consistency: Because both models use 1536 dimensions, you don't need to alter the VECTOR(1536) column type immediately.

Search Stability: Creating the index allows the engine to optimize the VECTOR_SEARCH function. As long as the incoming query vector (generated by the app) matches the dimensions of the stored vectors, the search will execute without a runtime dimension mismatch error.

Reference:

<https://docs.couchbase.com/cloud/n1ql/n1ql-language-reference/vectorfun.html>

42. Frage

.....

Sorgen Sie noch darum, dass Sie die Microsoft DP-800 Zertifizierungsprüfung nicht bestehen können? Dann sollen Sie sich an ITZert wenden. Wir können Sie die Top-Fähigkeit in der IT-Branche mitbringen, mit der Sie die Microsoft DP-800 Prüfung mühelos bestehen. Nach langjährigen Bemühungen beträgt die Bestehensrate bereits 100%. Wählen Sie ITZert, dann wählen Sie einen Weg zur glänzenden Zukunft.

DP-800 Vorbereitungsfragen: https://www.itzert.com/DP-800_valid-braindumps.html

Es gibt mehrere zugehörige Fragen und Antworten, mit denen Sie einen generalen Überblick und Eindruck über die Studienmaterialien von der DP-800 Vorbereitungsfragen - Developing AI-Enabled Database Solutions Prüfung haben können, Mit hochqualitativen und zuverlässigen DP-800 Studienführungsunterlagen versprechen wir Ihnen, dass Sie alle DP-800 Prüfungen ganz mühelos bestehen können, Wenn Sie sich an der Microsoft DP-800 Zertifizierungsprüfung beteiligen, wenden Sie sich ITZert an.

Der Wildling spuckte aus, Hogwarts sollte Dumbledores DP-800 Online Prüfung letzte Ruhestätte sein sagte Professor Flitwick, Es gibt mehrere zugehörige Fragen und Antworten, mit denen Sie einen generalen DP-800 Überblick und Eindruck über die Studienmaterialien von der Developing AI-Enabled Database Solutions Prüfung haben können.

Wir machen DP-800 leichter zu bestehen!

Mit hochqualitativen und zuverlässigen DP-800 Studienführungsunterlagen versprechen wir Ihnen, dass Sie alle DP-800 Prüfungen ganz mühelos bestehen können.

Wenn Sie sich an der Microsoft DP-800 Zertifizierungsprüfung beteiligen, wenden Sie sich ITZert an, Probieren Sie die Demo aus, um einen allgemeinen Überblick über unsere Produkte zu kriegen.

Sie können Praxis-Test VCE als Ihre eigenen Prüfung Simulation ablegen.

- Valid DP-800 exam materials offer you accurate preparation dumps Öffnen Sie www.echfrage.top geben Sie DP-800 ein und erhalten Sie den kostenlosen Download DP-800 Praxisprüfung
- DP-800 Unterlage DP-800 Online Praxisprüfung DP-800 Zertifikatsdemo Öffnen Sie die Website [www.itzert.com] Suchen Sie (DP-800) Kostenloser Download DP-800 Probesfragen
- DP-800 Braindumpsit Dumps PDF - Microsoft DP-800 Braindumpsit IT-Zertifizierung - Testking Examen Dumps Suchen Sie jetzt auf de.fast2test.com nach DP-800 um den kostenlosen Download zu erhalten DP-800 Praxisprüfung
- DP-800 Bestehen Sie Developing AI-Enabled Database Solutions! - mit höhere Effizienz und weniger Mühen Öffnen Sie die Webseite www.itzert.com und suchen Sie nach kostenloser Download von [DP-800] DP-800 Buch
- DP-800 Praxisprüfung DP-800 Simulationsfragen DP-800 Originale Fragen Öffnen Sie www.echfrage.top

- geben Sie ➔ DP-800 □ ein und erhalten Sie den kostenlosen Download □ DP-800 PDF Demo
- Valid DP-800 exam materials offer you accurate preparation dumps □ Suchen Sie einfach auf □ www.itzert.com □ nach kostenloser Download von ➤ DP-800 □ □ DP-800 Fragen Antworten
- DP-800 Zertifizierungsantworten □ DP-800 Prüfungsfrage □ DP-800 Praxisprüfung □ Geben Sie ⇒ www.zertpruefung.ch ⇐ ein und suchen Sie nach kostenloser Download von ➔ DP-800 □ □ DP-800 Fragen Antworten
- DP-800 Exam Fragen □ DP-800 Zertifizierungsantworten □ DP-800 Praxisprüfung □ Erhalten Sie den kostenlosen Download von [DP-800] mühelos über □ www.itzert.com □ *DP-800 Originale Fragen
- DP-800 Unterlage □ DP-800 Deutsch Prüfungsfragen □ DP-800 PDF Demo □ Suchen Sie auf ➔ www.zertpruefung.ch □ □ □ nach 【 DP-800 】 und erhalten Sie den kostenlosen Download mühelos □ DP-800 Tests
- DP-800 Zertifikatsdemo □ DP-800 Deutsch Prüfungsfragen □ DP-800 Originale Fragen □ [www.itzert.com] ist die beste Webseite um den kostenlosen Download von □ DP-800 □ zu erhalten □ DP-800 Deutsch Prüfungsfragen
- DP-800 Schulungsangebot, DP-800 Testing Engine, Developing AI-Enabled Database Solutions Trainingsunterlagen □ Suchen Sie jetzt auf ➔ www.zertsoft.com □ □ □ nach ➔ DP-800 □ □ □ um den kostenlosen Download zu erhalten □ DP-800 Fragen Und Antworten
- datatechcareers.com, www.stes.tyc.edu.tw, jakubehy1234306.wikirecognition.com, bookmarktiger.com, www.stes.tyc.edu.tw, techlearnersacademy.com, lexixgfb579891.blogspot.com, fanniegmdx058893.laowaiblog.com, jasonzbws518264.webdesign96.com, www.stes.tyc.edu.tw, Disposable vapes