

# DP-800 Exam Introduction - DP-800 Guide Torrent



What's more, part of that DumpStillValid DP-800 dumps now are free: <https://drive.google.com/open?id=1x4t3ufCAkQkWAeHq9EKDmsctiBXNv4Rv>

Our DP-800 study materials' developers to stand in the perspective of candidate, fully consider their material basis and actual levels of knowledge, formulated a series of scientific and reasonable learning mode, meet the conditions for each user to tailor their learning materials. What's more, our DP-800 Study Materials are cheap and cheap, and we buy more and deliver more. The more customers we buy, the bigger the discount will be. In order to make the user a better experience to the superiority of our DP-800 study materials.

## Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>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</li><li>CD pipelines using SQL Database Projects, and integrating SQL solutions with Azure services including Data API builder and monitoring tools.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>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.</li></ul>
Topic 3	<ul style="list-style-type: none"><li>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.</li></ul>

>> DP-800 Exam Introduction <<

## DP-800 Guide Torrent | Valid DP-800 Exam Vce

Once we have bought a practice materials, we may worry about that the version we bought cannot meet the need for the exam, so that we cannot know the latest information for the exam, if you worry about the questions like this and intend to join the DP-800 exam, just select the product of our company, because our products offer 365 days free update, it can help you to know about the latest information of the DP-800 Exam, so that you can change you strategies for the exam, besides downloading link of the update version will be sent to your email automatically by our systems. Using this, you can prepare for your test with ease.

## Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q20-Q25):

NEW QUESTION # 20

You have an Azure SQL database that contains the following tables and columns.

Embeddings in the NotesEmbeddings and DescriptionEmbeddings tables have been generated from values in the Description and notes columns of the Articles table by using different chunk sizes.

You need to perform approximate nearest neighbor (ANN) queries across both embedding tables. The solution must minimize the impact of using different chunk sizes.

What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer:**

Explanation:

Explanation:

The correct function is VECTOR\_SEARCH because the requirement is to perform approximate nearest neighbor (ANN) queries. Microsoft's SQL documentation states that VECTOR\_SEARCH is the function used for vector similarity search, and that an ANN index is used only with VECTOR\_SEARCH when a compatible vector index exists on the target column. By contrast, VECTOR\_DISTANCE calculates an exact distance and does not use a vector index for ANN retrieval.

The correct distance metric is cosine distance. Microsoft documents that VECTOR\_SEARCH supports cosine, dot, and euclidean metrics, and Microsoft guidance specifically notes that cosine similarity is commonly used for text embeddings. It also states that retrieval of the most similar texts to a given text typically functions better with cosine similarity, and that Azure OpenAI embeddings rely on cosine similarity to compute similarity between a query and documents. Since both NotesEmbeddings and DescriptionEmbeddings are text-derived embeddings and the goal is to minimize the impact of different chunk sizes, cosine is the best choice because it compares direction/angle rather than being as sensitive to vector magnitude as Euclidean distance.

**NEW QUESTION # 21**

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\_GENERME\_EMBEDOINGS 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. Normalize the vector lengths before storing new embeddings.
- **B. Regenerate embeddings for the existing rows.**
- C. Convert the Embedding column to nvarchar(max).
- D. Create a vector index on dbo.Products.Embedding.

**Answer: B**

Explanation:

When you change embedding models, the stored vectors should be treated as belonging to a different embedding space unless you intentionally keep the entire corpus consistent. Microsoft's vector guidance notes that when most or all embeddings are replaced with fresh embeddings from a new model, the recommended practice is to reload the new embeddings and, for large-scale replacement scenarios, consider dropping and recreating the vector index afterward so search quality remains predictable.

This question also says applications must continue to use VECTOR\_SEARCH without runtime errors .

VECTOR\_SEARCH requires compatible vector dimensions, and the vector column already exists. Azure OpenAI documentation shows that text-embedding-ada-002 is fixed at 1536 dimensions and text- embedding-3-small supports up to 1536 dimensions .

That means the migration can remain compatible with a VECTOR(1536) column, but the right implementation step is still to re-embed the existing rows so the table does not contain a mixed corpus produced by different models.

The other options are not appropriate:

\* B normalization does not solve a model migration problem.

\* C converting the vector column to nvarchar(max) would break vector-native search design.

\* D a vector index improves performance, but it does not migrate old embeddings to the new model.

**NEW QUESTION # 22**

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

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

**Answer: C**

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.

### NEW QUESTION # 23

You have an Azure SQL table that contains the following data.

You need to retrieve data to be used as context for a large language model (LLM). The solution must minimize token usage. Which format should you use to send the data to the LLM?

- A.
- B.
- C.
- D.

**Answer: D**

Explanation:

The correct choice is Option A because it provides the relevant semantic context the LLM needs while avoiding an unnecessary field that would add tokens without improving answer quality.

For LLM grounding and RAG-style context, Microsoft guidance emphasizes mapping and sending the fields that contain text pertinent to the use case. In this FAQ scenario, the useful context is the ProductName, the Question, and the Answer. Those three fields help the model understand both the subject domain and the actual Q & A pair. By contrast, FaqId is just a technical identifier and generally adds no semantic value for response generation, so including it wastes tokens.

That is why Option A is better than the others:

- \* Option A keeps the meaningful text fields and removes the low-value identifier.
- \* Option B is too minimal because it includes only the answer text as Prompt, which strips away the product and question context the LLM may need for accurate grounding.
- \* Option C keeps FaqId but omits ProductName, which can be important disambiguating context.
- \* Option D includes everything, but that does not minimize token usage because it keeps the unnecessary FaqId.

### NEW QUESTION # 24

You have a SQL database in Microsoft Fabric that contains a table named WebSite.Logs. WebSite.Logs stores application telemetry data. Website.Logs contains a nvarchar(max) column named log that stores JSON documents. You have a daily report that filters by the \$.severity JSON property and returns LogId, LogDateTime, and log.

The report frequently causes full table scans.

You need to modify Website.Logs to support efficient filtering by \$.severity and avoid key lookups for the columns returned by the

report.

How should you complete the Transact-SQL code to avoid full table scans? 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.

**Answer:**

Explanation:

Explanation:

The correct way to avoid full table scans here is to add a computed column that extracts the JSON scalar property with `JSON_VALUE`, and then create a nonclustered index on that computed column with the report's returned columns in the `INCLUDE` list. Microsoft's JSON indexing guidance specifically recommends creating a computed column that exposes the JSON property you filter on, using the same expression as in the query, and then indexing that computed column.

So the computed column must be:

```
AS JSON_VALUE([log], '$.severity') PERSISTED
```

This is correct because `$.severity` is a scalar JSON value, so `JSON_VALUE` is the proper function.

`JSON_QUERY` would be for extracting an object or array, not a scalar property. Microsoft also notes that persisted computed columns can improve access speed for JSON-derived values.

The index should then include:

```
INCLUDE (LogId, LogDateTime, [log])
```

That is the right covering strategy because the report filters by severity but returns `LogId`, `LogDateTime`, and `log`. Microsoft's guidance on included columns explains that nonkey included columns let a nonclustered index cover more queries and reduce extra lookups to the base table.

So the completed code is:

```
ALTER TABLE WebSite.Logs
```

```
ADD severity AS JSON_VALUE([log], '$.severity') PERSISTED;
```

```
GO
```

```
CREATE INDEX ix_severity
```

```
ON WebSite.Logs(severity)
```

```
INCLUDE (LogId, LogDateTime, [log]);
```

```
GO
```

**NEW QUESTION # 25**

.....

In the present market you are hard to buy the valid DP-800 study materials which are used to prepare the DP-800 exam like our DP-800 latest question. Both for the popularity in the domestic and the international market and for the quality itself, other kinds of study materials are incomparable with our DP-800 Test Guide and far inferior to them. Our DP-800 certification tool has their own fixed clients base in the domestic market and have an important share in the international market to attract more and more foreign clients.

**DP-800 Guide Torrent:** <https://www.dumpstillvalid.com/DP-800-prep4sure-review.html>

- DP-800 Exam Introduction Exam Instant Download | Updated Microsoft DP-800: Developing AI-Enabled Database Solutions  Search for ➡ DP-800  and download it for free on 「 [www.vce4dumps.com](http://www.vce4dumps.com) 」 website  Exam DP-800 Topic
- Pass your DP-800 exam in 2026 Smoothly!  Easily obtain ➡ DP-800  for free download through ➡ [www.pdfvce.com](http://www.pdfvce.com)   DP-800 Frenquent Update
- DP-800 Exam Introduction | High Pass-Rate Developing AI-Enabled Database Solutions 100% Free Guide Torrent  Go to website [ [www.examcollectionpass.com](http://www.examcollectionpass.com) ] open and search for ▶ DP-800 ◀ to download for free \* DP-800 Exam Cram Pdf
- Reliable DP-800 Exam Tips  DP-800 Frenquent Update  DP-800 Authorized Exam Dumps  Open website ▶ [www.pdfvce.com](http://www.pdfvce.com) ◀ and search for ➡ DP-800  for free download  DP-800 Valid Exam Vce
- 100% Pass Quiz 2026 Microsoft DP-800: Updated Developing AI-Enabled Database Solutions Exam Introduction  Go to website 「 [www.examcollectionpass.com](http://www.examcollectionpass.com) 」 open and search for ✓ DP-800  ✓  to download for free  Reliable DP-800 Exam Tips
- 100% Pass Quiz 2026 Microsoft DP-800: Updated Developing AI-Enabled Database Solutions Exam Introduction  Open  [www.pdfvce.com](http://www.pdfvce.com)  enter  DP-800  and obtain a free download ➡  Exam DP-800 Topic
- Reliable DP-800 Exam Tips  Latest DP-800 Dumps Sheet  DP-800 Authorized Exam Dumps  Immediately open

- ➡ [www.prep4sures.top](http://www.prep4sures.top) and search for “ DP-800 ” to obtain a free download Exam DP-800 Topic
- DP-800 Exam Cram Pdf DP-800 Frenquent Update DP-800 Frenquent Update Simply search for { DP-800 } for free download on [ [www.pdfvce.com](http://www.pdfvce.com) ] Exam DP-800 Topic
- DP-800 Reliable Study Guide Valid DP-800 Learning Materials Free Sample DP-800 Questions The page for free download of DP-800 on [www.testkingpass.com](http://www.testkingpass.com) will open immediately Exam DP-800 Topic
- Microsoft DP-800 Exam Introduction: Developing AI-Enabled Database Solutions - Pdfvce PDF Download Free (M) Download DP-800 for free by simply searching on [www.pdfvce.com](http://www.pdfvce.com) DP-800 Valid Exam Vce
- DP-800 Frenquent Update DP-800 Real Torrent Free Sample DP-800 Questions Easily obtain free download of { DP-800 } by searching on [www.exam4labs.com](http://www.exam4labs.com) DP-800 Authorized Exam Dumps
- [vinnytxrv075287 levitra-wiki.com](http://vinnytxrv075287 levitra-wiki.com), [carabbzz045406.blog4youth.com](http://carabbzz045406.blog4youth.com), [denisjejc872374.blogrelation.com](http://denisjejc872374.blogrelation.com), [haimamakp465006.izrablog.com](http://haimamakp465006.izrablog.com), [brianhwld962618.izrablog.com](http://brianhwld962618.izrablog.com), [brianaifu659286.wannawiki.com](http://brianaifu659286.wannawiki.com), [maerrbr355328.p2blogs.com](http://maerrbr355328.p2blogs.com), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [zoebisyw721107.wikisona.com](http://zoebisyw721107.wikisona.com), 119.29.134.108, Disposable vapes

BONUS!!! Download part of DumpStillValid DP-800 dumps for free: <https://drive.google.com/open?id=1x4t3ufCAkQkWAeHq9EKDmsctiBXNv4Rv>