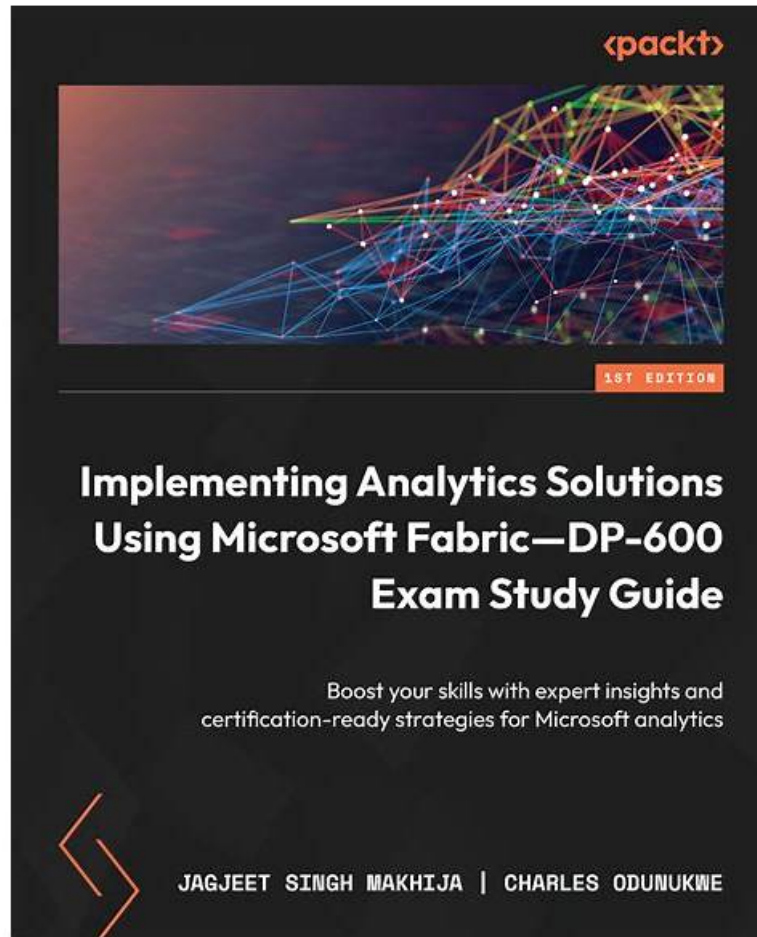


# Study Your Microsoft DP-800 Exam with Pass-Sure DP-800 Exam Discount: Developing AI-Enabled Database Solutions Efficiently



DOWNLOAD the newest BraindumpStudy DP-800 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1r8W3e7zsb5Bb7C0CCjbG3JfRgc5DWTBI>

We strongly recommend using our Developing AI-Enabled Database Solutions (DP-800) exam dumps to prepare for the Microsoft DP-800 certification. It is the best way to ensure success. With our Microsoft DP-800 practice questions, you can get the most out of your studying and maximize your chances of passing your Microsoft DP-800 Exam. BraindumpStudy Microsoft DP-800 practice test BraindumpStudy is the answer if you want to score higher in the DP-800 exam and achieve your academic goals.

## Microsoft DP-800 Exam Syllabus Topics:

Topic	Details
Topic 1	<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>
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>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>
---------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

>> DP-800 Exam Discount <<

## New DP-800 Exam Discount | Reliable Microsoft DP-800: Developing AI-Enabled Database Solutions 100% Pass

DP-800 latest study guide is the trustworthy source which can contribute to your actual exam test. If you are not sure about to pass your exam, you can rely on the DP-800 practice test for 100% pass. Microsoft DP-800 free pdf cram simulate the actual test, with the study of it, you can get a general understanding at first. After further practice with BraindumpStudy DP-800 Original Questions, you will acquire the main knowledge which may be tested in the actual test. At last, a good score is a little case.

### Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q19-Q24):

#### NEW QUESTION # 19

You need to recommend a solution to resolve the slow dashboard query issue. What should you recommend?

- A. On Lastupdatedutc, create a nonclustered index that includes Fleetid.
- B. On Fleetid, create a filtered index where lastupdatedutc > DATEADD(DAV, -7, SYSUTCOATETIME()).
- C. Create a clustered index on Lastupdatedutc.
- D. On Fleetid, create a nonclustered index that includes Lastupdatedutc, inginestatus, and BatteryHealth.

**Answer: D**

Explanation:

The best recommendation is B because the slow query filters on FleetId and returns LastUpdatedUtc, EngineStatus, and BatteryHealth. A nonclustered index with FleetId as the key column allows the optimizer to perform an index seek instead of a clustered index scan, and including the other selected columns makes the index covering, which reduces extra lookups and I/O. Microsoft's SQL Server indexing guidance states that a nonclustered index with included columns can significantly improve performance when all query columns are available in the index, because the optimizer can satisfy the query directly from the index. The query is:

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

Among the given choices, FleetId is the most important search argument because it appears in the WHERE predicate. Microsoft's index design guidance recommends putting columns used for searching in the key and using nonkey included columns to cover the rest of the query efficiently.

Why the other options are weaker:

\* A is not appropriate because changing the clustered index to LastUpdatedUtc would not target the main filter predicate on FleetId, and a table can have only one clustered index.

\* C makes LastUpdatedUtc the key, which is poor for a query whose primary filter is FleetId.

\* D is not the right answer here because the query requirement does not specify only recent rows, and filtered indexes are meant for a well-defined subset; this option also uses a time-based expression that is not aligned to the stated query pattern.

Strictly speaking, the most optimal design for both filtering and ordering would usually be a composite key like (FleetId, LastUpdatedUtc), but since that is not one of the available options, B is the correct exam answer.

#### NEW QUESTION # 20

You have an Azure SQL database that supports an OLTP application.

You need to write Transact-SQL code that returns blocking chain details. The output must return only sessions that are blocked or are blocking other sessions.

How should you complete the 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.

**Answer:**

Explanation:

Explanation:

\* CTE inner source # FROM sys.dm\_exec\_requests

\* Join after sys.dm\_exec\_sessions AS s # LEFT OUTER JOIN sys.dm\_exec\_requests

\* Text retrieval # OUTER APPLY sys.dm\_exec\_sql\_text(r.sql\_handle)

\* Input buffer retrieval # OUTER APPLY sys.dm\_exec\_input\_buffer(r.session\_id, r.request\_id) The correct drag-and-drop choices are based on how blocking-chain details are normally assembled in Azure SQL Database.

The CTE must read from sys.dm\_exec\_requests because the alias er is used with er.session\_id and er.

blocking\_session\_id, and those columns come from sys.dm\_exec\_requests. Microsoft documents that sys.

dm\_exec\_requests returns information about executing requests and includes the blocking\_session\_id column used to identify blockers.

After FROM sys.dm\_exec\_sessions AS s, the correct join is LEFT OUTER JOIN sys.dm\_exec\_requests so the query can still return sessions from sys.dm\_exec\_sessions even when a current request row is missing.

This is useful when showing sessions that are blocked or blocking, while still attempting to attach current request details when available.

For batch text, use OUTER APPLY sys.dm\_exec\_sql\_text(r.sql\_handle) because Microsoft documents sys.

dm\_exec\_sql\_text(sql\_handle) as the function that returns the SQL batch text for the specified sql\_handle.

For the input buffer, use OUTER APPLY sys.dm\_exec\_input\_buffer(r.session\_id, r.request\_id) because Microsoft documents that

sys.dm\_exec\_input\_buffer takes session\_id and request\_id and returns event\_info, which is commonly used when

sys.dm\_exec\_sql\_text is null or when you want the last command text.

So the completed code uses:

\* FROM sys.dm\_exec\_requests

\* LEFT OUTER JOIN sys.dm\_exec\_requests

\* OUTER APPLY sys.dm\_exec\_sql\_text(r.sql\_handle)

\* OUTER APPLY sys.dm\_exec\_input\_buffer(r.session\_id, r.request\_id)

**NEW QUESTION # 21**

Your team is developing an Azure SQL dataset solution from a locally cloned GitHub repository by using Microsoft Visual Studio Code and GitHub Copilot Chat.

You need to disable the GitHub Copilot repository-level instructions for yourself without affecting other users.

What should you do?

- A. From Visual Studio Code, modify your GitHub Copilot Chat user settings.
- B. Add a - debug flag when you start the GitHub Copilot Chat extension.
- C. Delete .github/copilot-instructions.md

**Answer: A**

Explanation:

GitHub documents that repository custom instructions for Copilot Chat can be disabled for your own use in the editor settings, and that doing so does not affect other users. In VS Code, this is controlled through settings related to instruction files, where you can disable the use of repository instruction files for your own environment.

The other options are incorrect:

\* B is not a documented mechanism for disabling repository-level Copilot instructions.

\* C would remove the repository instruction file itself and therefore affect everyone using that repository, which violates the requirement.

**NEW QUESTION # 22**

Your development team uses GitHub Copilot Chat in Microsoft SQL Server Management Studio (SSMS) to generate and run Transact-SQL queries against an Azure SQL database named DB1. DB1 contains tables that store sensitive customer data.

You need to ensure that any Transact SQL queries that run from GitHub Copilot Chat in SSMS are restricted by the same permissions as the developer's database login.

What prevents the GitHub Copilot Chat-run queries from accessing data beyond the developer ' s access?

- A. GitHub Copilot Chat runs queries by using the developer ' s database identity and permissions.
- B. GitHub Copilot Chat filters query results on (he client side to remove rows the developer is unauthorized to see.
- C. GitHub Copilot Chat runs queries in a read-only sandbox that is isolated from production database permissions.
- D. GitHub Copilot Chat uses different row-level security (RLS) policies than the developer.

**Answer: A**

Explanation:

The correct answer is B . Microsoft's SSMS Copilot documentation states that queries from Copilot in SSMS are executed under the context of the user's login and permissions , and that there are no separate permissions for Copilot in SSMS . That means Copilot-run Transact-SQL cannot access more data than the developer's own database principal is already allowed to access.

That is why the other options are incorrect:

\* A is wrong because Copilot does not use a separate read-only sandbox in place of database permissions.

\* C is wrong because enforcement is not a client-side filtering trick; it is enforced by the database security context of the current login.

\* D is wrong because Copilot does not apply a different RLS model from the developer; it simply runs under the same login context. So the security boundary is the developer's existing database identity and permissions .

### NEW QUESTION # 23

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. Add the schema to a GitHub Copilot instruction file.
- D. Use an MCP server

**Answer: D**

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 # 24

.....

One year of free Microsoft DP-800 test questions updates are included in the SnowPro Core Certification test DP-800 quiz package. This means that if any changes are made to the Developing AI-Enabled Database Solutions (DP-800) exam, you will be able to obtain the updated Microsoft DP-800 Test Questions preparation immediately. This is a great method to keep up to date on the latest Developing AI-Enabled Database Solutions (DP-800) questions information and ensure you pass the Developing AI-Enabled Database Solutions (DP-800) with ease.

**New DP-800 Dumps Files:** [https://www.braindumpstudy.com/DP-800\\_braindumps.html](https://www.braindumpstudy.com/DP-800_braindumps.html)

