DP-800 Online Bootcamps - DP-800 Valid Dumps Book

We provide free update to the client and the discounts to the old client. We provide free update of our DP-800 exam materials within one year and after one year the client can enjoy the 50% discounts. The old clients enjoy some certain discounts when they buy our DP-800 exam torrent. Our experts check whether there is the update of the test bank every day and if there is the system will send to the client automatically. We choose the most typical questions and answers which seize the focus and important information and the questions and answers are based on the real exam. So you can master the most important DP-800 Exam Torrent in the shortest time and finally pass the exam successfully.

Getting a certificate is not an easy thing for some of the candidates. DP-800 test dumps not only contain the quality, but also contain certain quality for your exam. Through using the DP-800 test dumps of us, you can pass the exam. In addition, DP-800 Test Dumps of us have the most of the knowledge points, and you can improve your ability in the process of learning. If you have any other questions about the DP-800 study materials, just contact us.

>> DP-800 Online Bootcamps <<

Microsoft DP-800 Valid Dumps Book | DP-800 Reliable Test Braindumps

ActualPDF DP-800 desktop and web-based practice exams are distinguished by their excellent features. The DP-800 web-based practice exam is supported by all operating systems and can be taken through popular browsers including Chrome, MS Edge, Internet Explorer, Opera, Firefox, and Safari. Windows computers can run the desktop Microsoft DP-800 Practice Test software. You won't require a live internet connection to use the desktop Microsoft exam simulation software once you've verified the product's license.

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	• 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 • CD pipelines using SQL Database Projects, and integrating SQL solutions with Azure services including Data API builder and monitoring tools.
Topic 3	• 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.

Microsoft Developing AI-Enabled Database Solutions Sample Questions (Q59-Q64):

NEW QUESTION # 59
You have a SQL database in Microsoft Fabric that contains a table named dbo.Orders, dbo.Orders has a clustered index, contains three years of data, and is partitioned by a column named OrderDate by month.
You need to remove all the rows for the oldest month. The solution must minimize the impact on other queries that access the data in dbo.orders.
Solution; Identify the partition scheme (or the oldest month, and then run the following Transact-SQL statement.
ALTER TABLE dbo.Orders
DROP PARTITION SCHEME (partition_scheme_name);
Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:
This also does not meet the goal. DROP PARTITION SCHEME removes the partition scheme object from the database; it is not the command used to remove just the rows for the oldest month from a partitioned table.
Microsoft's DROP PARTITION SCHEME documentation is explicit that the statement removes the partition scheme itself.
For removing only the oldest month's rows with minimal impact, Microsoft points to partition-level maintenance operations such as truncating a single partition on a partitioned table. That targets only the needed data subset and is more efficient for retention workloads.

NEW QUESTION # 60
Hotspot Question
You have an Azure SQL database that contains a table named knowledge_base.
knowledge_base stores human resources (HR) policy documents and contains columns named title, content, category, and embedding.
You have an application named App1. App1 queries two relational tables named employee_profiles and benefits_enrollment that contain HR data. App1 hosts a chatbot that calls a large language model (LLM) directly.
Users report that the chatbot answers general HR questions correctly but provides outdated or incorrect answers when policies change. The chatbot also fails to answer questions that reference internal policy documents by title or category.
You need to recommend a Retrieval Augmented Generation (RAG) solution to resolve the chatbot issues.
What should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

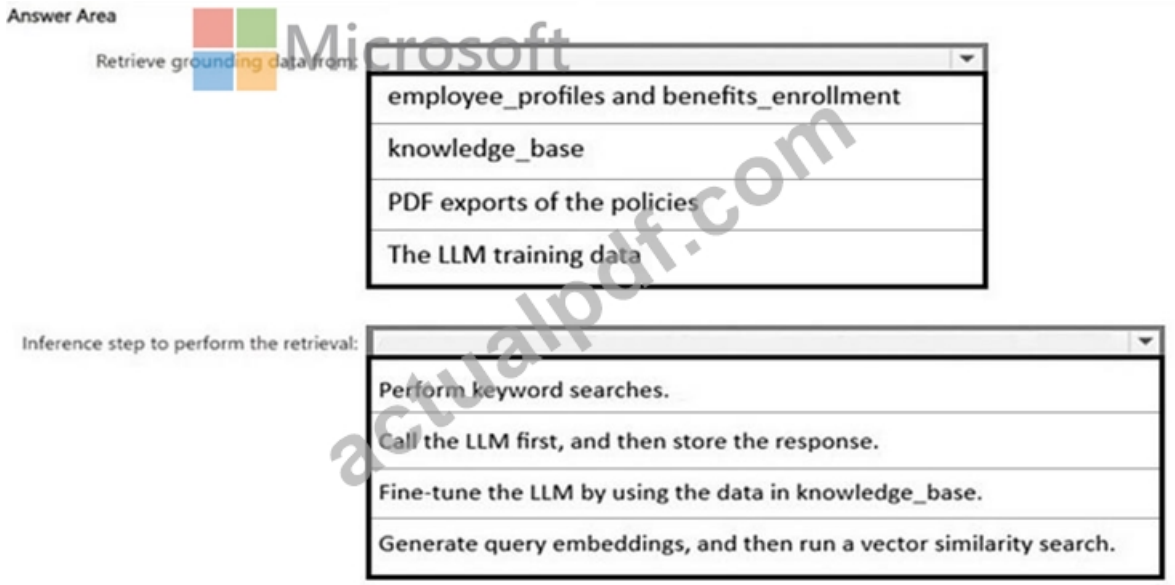

Answer:

Explanation:

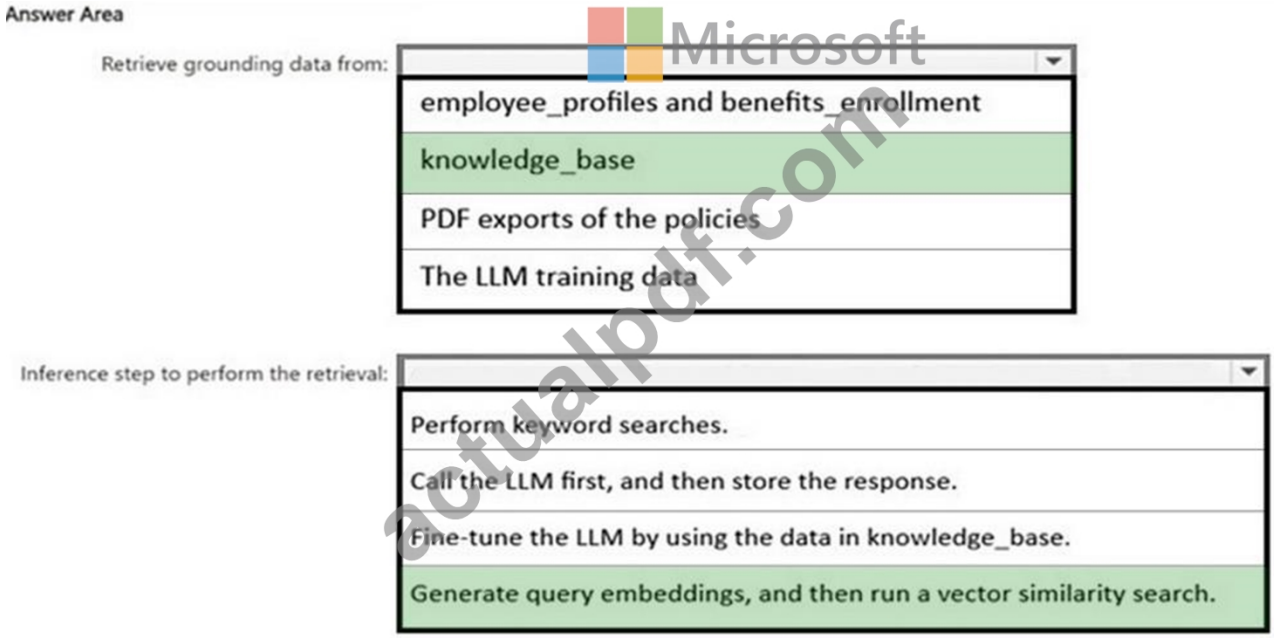

NEW QUESTION # 61
Case Study 2 - Fabrikam
Existing Environment
Azure Environment
Fabrikam has a single Azure subscription in the East US 2 Azure region. The subscription contains an Azure SQL database named DB1. DB1 contains the following tables:
* Patients
* Employees
* Procedures
* Transactions
* UsefulPrompts
* ProcedureDocuments
You store a column master key as a secret in Azure Key Vault.
You have an on-premises application named TransactionProcessing that uses a hard-coded username and password in a connection string to access DB1.
Problem Statements

Users report that after executing a long-running stored procedure named sp_UpdateProcedureForPatient, updates to the underlying data are sometimes inconsistent.
Requirements
Planned Changes
Fabrikam plans to manage all changes to Azure SQL Database objects by using source control in GitHub. Every pull request submitted to production will be validated before it can be merged.
Deployments must use the Release configuration.
Security Requirements
Fabrikam identifies the following security requirements:
* The TransactionProcessing application must use a passwordless connection to DB1.
* The Employees table contains two columns named TaxID and Salary that must be encrypted at rest.
* Auditors must have a tamper-evident history of transactions with cryptographic proof of changes to the employee data.
Database Performance Requirements
Records accessed by using sp_UpdateProcedureForPatient must NOT be changed by other transactions while the stored procedure runs.
AI Search, Embeddings, and Vector Indexing
Fabrikam identifies the following AI-related requirements:
* Queries to the ProcedureDocuments table must use Reciprocal Rank Fusion (RRF).
* Users must be able to query the data in DB1 by using prompts in Copilot in Microsoft Fabric.
* The UsefulPrompts table will store prompts that doctors can use to help diagnose patient illness by connecting to an Azure OpenAI endpoint.
Development Requirements
Fabrikam identifies the following development requirements:
* Provide the functionality to retrieve all the transactions of a given patient between two dates, showing a running total.
* Expose a Data API builder (DAB) configuration file to enable Azure services to perform the following operations over a REST API:
- Read data from the procedures table without authentication.
- Read and insert data into the Transactions table once authenticated.
- Execute the sp_UpdateProcedurePatient stored procedure.
* Provide the functionality to retrieve a list of the names of patients who underwent medical procedures during the last 30 days.
* Information for each medical procedure will be stored in a table. The table will be used with a large language model (LLM) for user querying and will have the following structure.

```
CREATE TABLE dbo.ProcedureDocuments
  (
    DocumentId INT IDENTITY PRIMARY KEY,
    SourceId NVARCHAR(200) NULL,
    Content NVARCHAR(MAX) NOT NULL,
    Embedding VECTOR(1536) NOT NULL,
    CreatedAt DATETIME2 NOT NULL DEFAULT SYSUTCDATETIME()
  );
```


DAB
You create a DAB configuration file that meets the development requirements for DB1 and includes the following entities.

```
"entities": {
  "Procedures": {
    "source": "dbo.Procedures",
    "rest": true,
    "graphql": true,
    "permissions": [
      {
        "role": "anonymous",
        "actions": [ "read" ]
   }
    ]
  },
  "Transactions": {
    "source": "dbo.Transactions",
    "rest": true,
    "graphql": true,
    "permissions": [
      {
        "role": "authenticated",
        "actions": [ "read", "create" ]
      }
    ]
  },
  "UpdateProcedurePatient": {
    "source": "dbo.sp_ UpdateProcedurePatient",
    "rest": {
      "enabled": true,
      "method": "post",
      "path": "/procedurepatient"
    },
    "graphql": false,
    "permissions": [
      {
        "role": "authenticated",
        "actions": [ "execute" ]
      }
    ]
  }
}
```


You need to use the UsefulPrompts table as defined in the AI requirements. Which stored procedure should you use?

- A. sp_addendpoint
- B. xp_cmdshell
- C. sp_OACreate
- D. sp_invoke_external_rest_endpoint

Answer: D

Explanation:
Scenario:
The UsefulPrompts table will store prompts that doctors can use to help diagnose patient illness by connecting to an Azure OpenAI endpoint.
The system stored procedure sp_invoke_external_rest_endpoint is used to connect an Azure SQL Database to an Azure OpenAI endpoint.This procedure allows you to call HTTPS REST endpoints directly from your database, enabling the integration of generative AI or embedding models into your SQL workflows without an intermediate application layer.

Reference:
https://blog.fabric.microsoft.com/en-gb/blog/ai-ready-apps-from-rag-to-chat-interacting-with-sql- database-in-microsoft-fabric-using-graphql-and-mcp

NEW QUESTION # 62
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a SQL database in Microsoft Fabric that contains a table named dbo.Orders.
dbo.Orders has a clustered index, contains three years of data, and is partitioned by a column named OrderDate by month.
You need to remove all the rows for the oldest month. The solution must minimize the impact on other queries that access the data in dbo.Orders.
Solution: Identify the partition scheme for the oldest month, and then run the following Transact- SQL statement.
ALTER TABLE dbo.Orders
DROP PARTITION SCHEME (partition_scheme_name);
Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:
Correct:
* Identify the partition number for the oldest month, and then run the following Transact-SQL statement.
TRUNCATE TABLE dbo.Orders
WITH (PARTITIONS (partition number));
The best Transact-SQL statement to remove all rows for the oldest month while minimizing the impact on other queries is TRUNCATE TABLE with a WITH (PARTITIONS (...)) clause.
Why TRUNCATE TABLE ... WITH (PARTITIONS (...)) is Best
Efficiency: TRUNCATE TABLE is a Data Definition Language (DDL) operation that removes data by deallocating the data pages, which is a metadata operation and is very fast, regardless of the amount of data in the partition.
Minimal Logging: It uses less transaction log space compared to a DELETE statement, which logs each row deletion individually.
Low Impact on Concurrency: It performs a quick, partition-specific operation. A row-by-row DELETE would be a long-running transaction and could cause locking and blocking issues for other queries accessing the table.
Data Integrity: Because the table has a clustered index and is partitioned by the same column (aligned indexes), the TRUNCATE PARTITION operation is a fast, partition-level maintenance operation that targets only that specific data subset.
Incorrect:
* : Identify the partition scheme for the oldest month, and then run the following Transact-SQL statement.
ALTER TABLE dbo.Orders
DROP PARTITION SCHEME (partition_scheme_name);
The DROP PARTITION SCHEME statement removes the partition scheme object from the database but does not remove the data itself or free up the space, and it requires all tables to be moved off the scheme first, which is a complex operation. This does not meet the goal of removing the data efficiently.
* Run the following Transact-SQL statement.
DELETE FROM dbo.Orders
WHERE OrderDate < DATEADD(month, -36, SYSUTCDATETIME());
A standard DELETE statement, even with a WHERE clause that uses the partition column, can be a time-consuming, logged operation that causes locking and blocking on the main table, negatively impacting performance.
Reference:
https://stackoverflow.com/questions/63632963/truncate-partition-vs-drop-partition-performace- wise-which-one-is-efficient-an

NEW QUESTION # 63
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.
```
SELECT p3.PersonId, p3.DisplayName FROM dbo.Person AS p1
JOIN dbo.Knows AS k1 ON 1 = 1
JOIN dbo.Person AS p2 ON 1 = 1
JOIN dbo.Knows AS k2 ON 1 = 1
JOIN dbo.Person AS p3 ON 1 = 1
WHERE p1.PersonId = @startPersonId
  AND MATCH(p3<-(k2)-p2<-(k1)-p1);
```
- B.
```
SELECT p3.PersonId, p3.DisplayName
FROM dbo.Person AS p1, dbo.Knows AS k1, dbo.Person AS p2, dbo.Knows AS k2, dbo.Person AS p3
WHERE p1.PersonId = @StartPersonId
  AND MATCH(p1-(k1)->p2-(k2)->p3);
```
- C.
```
SELECT p2.PersonId, @StartPersonId
FROM dbo.Person AS p1, dbo.Knows AS k1, dbo.Person AS p2, dbo.Knows AS k2, dbo.Person AS p3
WHERE p1.DisplayName = p2.DisplayName
  AND MATCH(p1-(k1)->p2-(k2)->p3);
```
- D.
```
SELECT p3.PersonId, p3.DisplayName
FROM dbo.Person AS p1, dbo.Knows AS k1, dbo.Person AS p2, dbo.Knows AS k2, dbo.Person AS p3
WHERE p1.PersonId = @StartPersonId
  AND MATCH(p1-(k1)->p2)  AND MATCH(p2-(k2)->p3);
```

Answer: B

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.

Table names	Column names
CustomerFeedback	• FeedbackId (int) (primarykey) • FeedbackJson (nvarchar (max))
Fleets	• FleetId (int) (primarykey) • FleetName (nvarchar(100)) • Description (nvarchar(256))
MaintenanceEvents	• MaintenanceId (int) (primarykey) • VehicleId (int) • LastModifiedUTC (datetime2) • Description (nvarchar(256))
SupportTickets	• TicketId (int) (primarykey) • FleetId (int) • CreatedUtc (datetime2)
UserAccounts	• UserId (int) (primarykey) • UserPrincipalName (nvarchar(256)) • JobRole (nvarchar(256)) • StartDate (datetime2)
VehicleHealthSummary	• IncidentId (int) (primarykey) • VehicleId (int) • 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.

Requirements

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 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 AI-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 # 64
.....

We are never complacent about our achievements, so all content of our DP-800 exam questions are strictly researched by proficient experts who absolutely in compliance with syllabus of this exam. Accompanied by tremendous and popular compliments around the world, to make your feel more comprehensible about the DP-800 study prep, all necessary questions of knowledge concerned with the exam are included into our DP-800 simulating exam.

DP-800 Valid Dumps Book: https://www.actualpdf.com/DP-800_exam-dumps.html

- Pass Guaranteed Quiz 2026 Updated Microsoft DP-800 Online Bootcamps □ Download □ DP-800 □ for free by simply entering 《 www.dumpsquestion.com 》 website □DP-800 Reliable Test Duration
- New DP-800 Test Prep □ Interactive DP-800 EBook □ Test DP-800 King □ The page for free download of " DP-800 " on " www.pdfvce.com " will open immediately □Reliable DP-800 Test Tips
- Latest DP-800 Exam Vce □ DP-800 Latest Exam Pass4sure □ Valid Braindumps DP-800 Pdf □ Search for □ DP-800 □ and download it for free immediately on □ www.exam4labs.com □ □DP-800 Latest Exam Pass4sure
- DP-800 Online Bootcamps - High-quality Microsoft DP-800 Valid Dumps Book: Developing AI-Enabled Database Solutions □ Search for 《 DP-800 》 and download exam materials for free through □ www.pdfvce.com □ □Detailed DP-800 Answers
- New DP-800 Test Labs □ Valuable DP-800 Feedback □ DP-800 Reliable Test Duration □ Search for ➠ DP-800 □ and download exam materials for free through [www.prep4sures.top] □Reliable DP-800 Test Tips
- Microsoft DP-800 Exam Dumps Are Available At A Cheap Price □ Download □ DP-800 □ for free by simply entering ▹ www.pdfvce.com ◃ website □Valuable DP-800 Feedback
- Valid Dumps DP-800 Free □ DP-800 Latest Exam Pass4sure □ Valid Braindumps DP-800 Pdf □ The page for free download of ➡ DP-800 □□□ on ➽ www.pdfdumps.com □ will open immediately □DP-800 Latest Test Braindumps
- Quiz 2026 High Hit-Rate Microsoft DP-800: Developing AI-Enabled Database Solutions Online Bootcamps □ Download ✔ DP-800 □✔□ for free by simply searching on ▸ www.pdfvce.com ◂ □Interactive DP-800 EBook
- 2026 Latest 100% Free DP-800 – 100% Free Online Bootcamps | DP-800 Valid Dumps Book □ Search for ➠ DP-800 □ on ⇛ www.prepawayete.com ⇚ immediately to obtain a free download ☢Test DP-800 King
- Interactive DP-800 EBook □ DP-800 Valid Braindumps Ppt ➽ Test DP-800 Cram □ The page for free download of { DP-800 } on □ www.pdfvce.com □ will open immediately □Valid Braindumps DP-800 Pdf
- DP-800 Latest Exam Pass4sure □ Valid DP-800 Exam Bootcamp □ Test DP-800 Cram □ Immediately open ☀ www.prepawayexam.com □☀□ and search for { DP-800 } to obtain a free download □Valid Braindumps DP-800 Pdf
- www.stes.tyc.edu.tw, neilduyi314352.bloggerchest.com, bookmarkingalpha.com, thebookmarkking.com, atozbookmark.com, rorybuym959779.ssnblog.com, hamzaagnc938269.wiki-jp.com, declanqzjh851151.wikifrontier.com, deannazaew891500.blogdun.com, bookmarkquotes.com, Disposable vapes