Free Questions for DP-300 by braindumpscollection

Shared by Mullen on 09-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: Hotspot

You need to recommend the appropriate purchasing model and deployment option for the 30 new databases. The solution must meet the technical requirements and the business requirements.

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

NOTE: Each correct selection is worth one point.

Image not found or type unknown

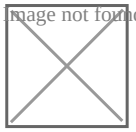

Explanation:

https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-pool-overview

https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview

Question 2

Question Type: MultipleChoice

Task 1

In an Azure SQL database named db1, you need to enable page compression on the PK_SalesOrderHeader_SalesOrderlD clustered index of the SalesLT.SalesOrderHeader table.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To enable page compression on the PK_SalesOrderHeader_SalesOrderlD clustered index of the SalesLT.SalesOrderHeader table in db1, you can use the following Transact-SQL script:

```
-- Connect to the Azure SQL database named db1

USE db1;
```

```
GO

-- Enable page compression on the clustered index

ALTER INDEX PK_SalesOrderHeader_SalesOrderlD ON SalesLT.SalesOrderHeader

REBUILD WITH (DATA_COMPRESSION = PAGE);

GO
```

This script will rebuild the clustered index with page compression, which can reduce the storage space and improve the query performance

The script solution consists of three parts:

The first part isUSE db1; GO. This part connects to the Azure SQL database named db1, where the SalesLT.SalesOrderHeader table is located. TheGOcommand separates the batches of Transact-SQL statements and sends them to the server.

The second part isALTER INDEX PK_SalesOrderHeader_SalesOrderlD ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO. This part enables page compression on the clustered index named PK_SalesOrderHeader_SalesOrderlD, which is defined on the SalesLT.SalesOrderHeader table. TheALTER INDEXstatement modifies the properties of an existing index. TheREBUILDoption rebuilds the index from scratch, which is required to change the compression setting. TheDATA_COMPRESSION = PAGEoption specifies that page compression is applied to the index, which means that both row and prefix compression are used. Page compression can reduce the storage space and improve the query performance by compressing the data at the page level. TheGOcommand ends the batch of statements.

The third part is optional, but it can be useful to verify the compression status of the index. It isSELECT name, index_id, data_compression_desc FROM sys.indexes WHERE object_id = OBJECT_ID('SalesLT.SalesOrderHeader');. This part queries the

sys.indexes catalog view, which contains information about the indexes in the database. TheSELECTstatement returns the name, index_id, and data_compression_desc columns for the indexes that belong to the SalesLT.SalesOrderHeader table. TheOBJECT_IDfunction returns the object identification number for the table name. The data_compression_desc column shows the compression type of the index, which should be PAGE for the clustered index after the script is executed.

These are the steps of the script solution for enabling page compression on the clustered index of the SalesLT.SalesOrderHeader table in db1.

Question 3

Question Type: **MultipleChoice**

Task 2

You need to configure your user account as the Azure AD admin for the server named sql3700689S.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To configure your user account as the Azure AD admin for the server named sql3700689S, you can use the Azure portal or the Azure CLI. Here are the steps for both methods:

Using the Azure portal:

Go to the Azure portal and selectSQL Server -- Azure Arc.

Select the server namedsql3700689Sand click onActive Directory admin.

Click onSet adminand choose your user account from the list of Azure AD users.

Click onSelectand thenSaveto confirm the change.

You can verify the Azure AD admin by clicking onActive Directory adminagain and checking the current admin.

Using the Azure CLI:

Install the Azure CLI and log in with your Azure account.

Run the following command to get the object ID of your user account:az ad user show --id <your-user-name> --query objectId -o tsv

Run the following command to set your user account as the Azure AD admin for the server:az sql server ad-admin create --server sql3700689S --object-id <your-object-id> --display-name <your-user-name>

You can verify the Azure AD admin by running the following command:az sql server ad-admin show --server sql3700689S

These are the steps to configure your user account as the Azure AD admin for the server named sql3700689S.

Question 4

Question Type: **MultipleChoice**

Task 3

You need to ensure that all queries executed against dbl are captured in the Query Store.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To ensure that all queries executed against dbl are captured in the Query Store, you need to enable the Query Store feature for the database and set the query capture mode to ALL.The Query Store feature provides you with insight on query plan choice and performance for Azure SQL Database1.The query capture mode controls whether all queries or only a subset of queries are tracked2.

Here are the steps to enable the Query Store and set the query capture mode to ALL for the database dbl:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server.

Select the database dbl and click onQuery Performance Insightin the left menu.

Click onConfigure Query Storeand turn on theQuery Storeswitch.

In theQuery Capture Modedropdown, selectAlland click onSave.

Using Transact-SQL statements:

Connect to the Azure SQL Database server and the database dbl using SQL Server Management Studio or Azure Data Studio.

Run the following command to enable the Query Store for the database:ALTER DATABASE dbl SET QUERY_STORE = ON;

Run the following command to set the query capture mode to ALL for the database:ALTER DATABASE dbl SET QUERY_STORE (QUERY_CAPTURE_MODE = ALL);

These are the steps to ensure that all queries executed against dbl are captured in the Query Store.

Question 5

Question Type: MultipleChoice

Task 5

You need to configure a disaster recovery solution for db1. When a failover occurs, the connection strings to the database must remain the same. The secondary server must be in the West US 3 Azure region.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To configure a disaster recovery solution for db1, you can use the failover groups feature of Azure SQL Database.Failover groups allow you to manage the replication and failover of a group of databases across different regions with the same connection strings1.You can also use active geo-replication as an alternative, but you will need to update the connection strings manually after a failover2.

Here are the steps to create a failover group for db1 with the secondary server in the West US 3 region:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

SelectFailover groupsin the left menu and click onAdd group.

Enter a name for the failover group and selectWest US 3as the secondary region.

Click onCreate a new serverand enter the details for the secondary server, such as server name, admin login, password, and subscription.

Click onSelect existing database(s)and choose db1 from the list of databases on the primary server.

Click onConfigure failover policyand select the failover mode, grace period, and read-write failover endpoint mode according to your preferences.

Click onCreateto create the failover group and start the replication of db1 to the secondary server.

Using PowerShell commands:

Install the Azure PowerShell module and log in with your Azure account.

Run the following command to create a new server in the West US 3 region:New-AzSqlServer -ResourceGroupName <your-resource-group-name> -ServerName <your-secondary-server-name> -Location 'West US 3' -SqlAdministratorCredentials $(New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList '<your-admin-login>', $(ConvertTo-SecureString -String '<your-password>' -AsPlainText -Force))

Run the following command to create a new failover group with db1:New-AzSqlDatabaseFailoverGroup -ResourceGroupName <your-resource-group-name> -ServerName <your-primary-server-name> -PartnerResourceGroupName <your-resource-group-name> -PartnerServerName <your-secondary-server-name> -FailoverGroupName <your-failover-group-name> -Database db1 -FailoverPolicy Manual -GracePeriodWithDataLossHours 1 -ReadWriteFailoverEndpoint 'Enabled'

You can modify the parameters of the command according to your preferences, such as the failover policy, grace period, and read-write failover endpoint mode.

These are the steps to create a failover group for db1 with the secondary server in the West US 3 region.

Question 6

Question Type: **MultipleChoice**

Task 4

You need to enable change data capture (CDC) for db1.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To enable change data capture (CDC) for db1, you need to run the stored procedure sys.sp_cdc_enable_db in the database context.CDC is a feature that records activity on a database when tables and rows have been modified1.CDC can be used for various scenarios, such as data synchronization, auditing, or ETL processes2.

Here are the steps to enable CDC for db1:

Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.

Open a new query window and run the following command:EXEC sys.sp_cdc_enable_db; GO

This command will enable CDC for the database and create the cdc schema, cdc user, metadata tables, and other system objects for the database3.

To verify that CDC is enabled for db1, you can query the is_cdc_enabled column in the sys.databases catalog view. The value should be 1 for db1.

These are the steps to enable CDC for db1

Question 7

Question Type: MultipleChoice

Task 7

You plan to create an automation runbook that will create database users in db1 from Azure AD identities. You need to configure sq1370O6895 to support the creation of new database users.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To configure sq1370O6895 to support the creation of new database users from Azure AD identities, you need to do the following steps:

Set up a Microsoft Entra tenant and associate it with your Azure subscription.You can use the Microsoft Entra portal or the Azure portal to create and manage your Microsoft Entra users and groups12.

Configure a Microsoft Entra admin for sq1370O6895.You can use the Azure portal or the Azure CLI to set a Microsoft Entra user as the admin for the server34.The Microsoft Entra admin can create other database users from Microsoft Entra identities5.

Connect to db1 using the Microsoft Entra admin account and run the following Transact-SQL statement to create a new database user from a Microsoft Entra identity:CREATE USER [Microsoft Entra user name] FROM EXTERNAL PROVIDER;6You can replace the Microsoft Entra user name with the name of the user or group that you want to create in the database.

Grant the appropriate permissions to the new database user by adding them to a database role or granting them specific privileges. For example, you can run the following Transact-SQL statement to add the new user to the db_datareader role:ALTER ROLE db_datareader ADD MEMBER [Microsoft Entra user name];

These are the steps to configure sq1370O6895 to support the creation of new database users from Azure AD identities.

Question 8

Question Type: **MultipleChoice**

Task 8

You plan to perform performance testing of db1.

You need prevent db1 from reverting to the last known good query plan.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To prevent db1 from reverting to the last known good query plan, you need to disable the automatic plan correction feature for the database.This feature is enabled by default and allows the Query Store to detect and fix plan performance regressions by forcing the last good plan1.However, if you want to test the performance of different plans without interference from the Query Store, you can turn off this feature by using the ALTER DATABASE SCOPED CONFIGURATION statement2.

Here are the steps to disable the automatic plan correction feature for db1:

Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.

Open a new query window and run the following command:ALTER DATABASE SCOPED CONFIGURATION SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = OFF); GO

This command will disable the automatic plan correction feature for db1 and allow the Query Optimizer to choose the best plan based on the current statistics and parameters3.

To verify that the automatic plan correction feature is disabled for db1, you can query the sys.database_scoped_configurations catalog view. The value of the force_last_good_plan column should be 0 for db1.

These are the steps to disable the automatic plan correction feature for db1.

Question 9

Question Type: MultipleChoice

Task 9

You need to generate an email alert to admin@contoso.com when CPU percentage utilization for db1 is higher than average.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To generate an email alert to admin@contoso.com when CPU percentage utilization for db1 is higher than average, you can use the Azure portal to create an alert rule based on the CPU percentage metric. Here are the steps to do that:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

SelectAlertsin the Monitoring section and click onNew alert rule.

In the Condition section, clickAddand select theCPU percentagemetric.

In the Configure signal logic page, set the threshold type toDynamic.This will compare the current metric value to the historical average and trigger the alert when it deviates significantly1.

Set the operator toGreater than, the aggregation type toAverage, the aggregation granularity to1 minute, and the frequency of evaluation to5 minutes.

ClickDoneto save the condition.

In the Action group section, clickCreateand enter a name and a short name for the action group.

In the Notifications section, clickAddand selectEmail/SMS message/Push/Voice.

Enteradmin@contoso.comin the Email field and clickOK.

ClickOKto save the action group.

In the Alert rule details section, enter a name and a description for the alert rule, choose a severity level, and make sure the rule is enabled.

ClickCreate alert ruleto create the alert rule.

This alert rule will send an email to admin@contoso.com when the CPU percentage utilization for db1 is higher than average.You can also add other actions to the alert rule, such as calling a webhook or running an automation script

Question 10

Question Type: MultipleChoice

Task 11

You have a legacy application written for Microsoft SQL Server 2012. The application will be the only application that accesses db1 You need to ensure that db1 is compatible with all the features and syntax of SQL Server 2012.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To ensure that db1 is compatible with all the features and syntax of SQL Server 2012, you need to set the compatibility level of the database to 110, which is the compatibility level for SQL Server 20121.The compatibility level affects the behavior of certain Transact-SQL statements and features, and determines how the database engine interprets the SQL code2.

You can set the compatibility level of db1 by using the Azure portal or Transact-SQL statements. Here are the steps for both methods:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

Select the database db1 and click onQuery Performance Insightin the left menu.

Click onConfigure Query Storeand select110from theCompatibility leveldropdown list.

Click onSaveto apply the change.

Using Transact-SQL statements:

Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.

Open a new query window and run the following command:ALTER DATABASE db1 SET COMPATIBILITY_LEVEL = 110; GO

This command will set the compatibility level of db1 to 110, which is equivalent to SQL Server 2012.

These are the steps to set the compatibility level of db1 to 110.

Question 11

Question Type: MultipleChoice

Task 10

You need to protect all the databases on sql37006S95 from SQL injection attacks.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

SQL injection attacks are a type of cyberattack that exploit a vulnerability in the application code that interacts with the database.An attacker can inject malicious SQL statements into the user input, such as a form field or a URL parameter, and execute them on the database server, resulting in data theft, corruption, or unauthorized access1.

To protect all the databases on sql37006S95 from SQL injection attacks, you need to follow some best practices for securing your application and database layers. Here are some of the recommended steps:

Use parameterized queries or stored procedures to separate the SQL code from the user input.This will prevent the user input from being interpreted as part of the SQL statement and avoid SQL injection23.

Validate and sanitize the user input before passing it to the database.This will ensure that the input conforms to the expected format and type, and remove any potentially harmful characters or keywords4.

Implement least privilege access for the database users and roles.This will limit the permissions and actions that the application can perform on the database, and reduce the impact of a successful SQL injection attack5.

Enable Advanced Threat Protection for Azure SQL Database. This is a feature that detects and alerts you of anomalous activities and potential threats on your database, such as SQL injection, brute force attacks, or unusual access patterns. You can configure the alert settings and notifications using the Azure portal or PowerShell.

These are some of the steps to protect all the databases on sql37006S95 from SQL injection attacks.

Question 12

Question Type: **MultipleChoice**

Task 12

You need to configure high availability for dbl. The solution must tolerate the loss of an Azure datacenter without data loss or the need to modify application connection strings.

Options:

A) See the explanation part for the complete Solution

Answer:

A

Explanation:

To configure high availability for dbl, you can use the failover groups feature of Azure SQL Database.Failover groups allow you to manage the replication and failover of a group of databases across different regions with the same connection strings1. You can choose all, or a subset of, user databases in a logical server to be replicated to another logical server in a different region. You can also specify the failover policy, such as manual or automatic, and the grace period for data loss.

Here are the steps to create a failover group for dbl:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts dbl.

SelectFailover groupsin the left menu and click onAdd group.

Enter a name for the failover group and select a secondary region that is different from the primary region.

Click onCreate a new serverand enter the details for the secondary server, such as server name, admin login, password, and subscription.

Click onSelect existing database(s)and choose dbl from the list of databases on the primary server.

Click onConfigure failover policyand select the failover mode, grace period, and read-write failover endpoint mode according to your preferences.

Click onCreateto create the failover group and start the replication of dbl to the secondary server.

Using PowerShell commands:

Install the Azure PowerShell module and log in with your Azure account.

Run the following command to create a new server in the secondary region:New-AzSqlServer -ResourceGroupName <your-resource-group-name> -ServerName <your-secondary-server-name> -Location '<secondary-region-name>' -SqlAdministratorCredentials $(New-Object -TypeName System.Management.Automation.PSCredential -ArgumentList '<your-admin-login>', $(ConvertTo-SecureString -String '<your-password>' -AsPlainText -Force))

Run the following command to create a new failover group with dbl:New-AzSqlDatabaseFailoverGroup -ResourceGroupName <your-resource-group-name> -ServerName <your-primary-server-name> -PartnerResourceGroupName <your-resource-group-name> -PartnerServerName <your-secondary-server-name> -FailoverGroupName <your-failover-group-name> -Database

dbl -FailoverPolicy <manual-or-automatic> -GracePeriodWithDataLossHours <grace-period-in-hours> -ReadWriteFailoverEndpoint '<enabled-or-disabled>'

You can modify the parameters of the command according to your preferences, such as the failover policy, grace period, and read-write failover endpoint mode.

These are the steps to create a failover group for dbl