



*Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.*



DP-300 Dumps  
DP-300 Braindumps  
DP-300 Real Questions  
DP-300 Practice Test  
DP-300 Actual Questions



**Microsoft**

# DP-300

*Administering Relational Databases on Microsoft Azure*



### Question: 87

You are designing a dimension table in an Azure Synapse Analytics dedicated SQL pool.

You need to create a surrogate key for the table. The solution must provide the fastest query performance.

What should you use for the surrogate key?

- A. an IDENTITY column
- B. a GUID column
- C. a sequence object

**Answer:** A

Explanation:

Dedicated SQL pool supports many, but not all, of the table features offered by other databases.

Surrogate keys are not supported. Implement it with an Identity column.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesoverview>

### Question: 88

Topic 2, Contoso Ltd

Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this

section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

## Overview

### Existing Environment

Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.

### Active Directory

Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active Directory.

### Database Environment

Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

| Name | Role                                      |
|------|---|
| SQL1 | Primary data warehouse                    |
| SQL2 | Secondary data warehouse                  |
| SQL3 | Extract, transform, and load (ETL) server |

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.

The on-premises datacenter contains a PostgreSQL server that has a 50-TB database.

### Current Business Model

Contoso uses Microsoft SQL Server Integration Services (SSIS) to create flat files for customers. The customers receive the files by using FTP.

### Requirements

### Planned Changes

Contoso plans to move to a model in which they deliver data to customer databases that run as platform as a service (PaaS) offerings. When a customer establishes a service agreement with Contoso, a separate resource group that contains an Azure SQL database will be provisioned for the customer. The database will have a complete copy of the financial data. The data to which each customer will have access will depend on the service agreement tier. The customers can change tiers by changing their service agreement.

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.

Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months.

Start onboarding customers to the new PaaS solution within six months.

## Business Goals

Contoso identifies the following business requirements:

Use built-in Azure features whenever possible.

Minimize development effort whenever possible.

Minimize the compute costs of the PaaS solutions.

Provide all the customers with their own copy of the database by using the PaaS solution. Provide the customers with different table and row access based on the customer's service agreement.

In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.

Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso.

## Technical Requirements

Contoso identifies the following technical requirements:

Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure AD of Contoso to minimize guest users.

All customers must have their own resource group, Azure SQL server, and Azure SQL database. The deployment of resources for each customer must be done in a consistent fashion.

Users must be able to review the queries issued against the PaaS databases and identify any new objects created.

Downtime during the PostgreSQL database migration must be minimized.

## Monitoring Requirements

Contoso identifies the following monitoring requirements:

Notify administrators when a PaaS database has a higher than average CPU usage.


Use a single dashboard to review security and audit data for all the PaaS databases.

Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.

Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues automatically whenever possible.

## PaaS Prototype

During prototyping of the PaaS solution in Azure, you record the compute utilization of a customer’s Azure SQL database as shown in the following exhibit.



Role Assignments

For each customer’s Azure SQL Database server, you plan to assign the roles shown in the following exhibit.

+

Add

Edit columns

Refresh

Remove

Got feedback?

Check access

Role assignments

Deny assignments

Classic administrators

Roles

Manage access to Azure resources for users, groups, service principals and managed identities at this scope by creating role assignments. [Learn more](#)

Number of role assignments for this subscription 15 2000

Name

Type

Role

Scope

Search by name or email

Groups

2 selected

All scopes

Group by

Role

Showing a filtered set of results. Total number of role assignments: 15

2 items (2 Groups)

Name

Type

Role

Scope

Contributor

DB

DBAGroup1

Group

Contributor

This resource

SQL DB Contributor

DB

DBAGroup2

Group

SQL DB Contributor

This resource

HOTSPOT

You are evaluating the role assignments.

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                    |
|---|-----------------------|-----------------------|
| DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.     | <input type="radio"/> | <input type="radio"/> |
| DBAGroup1 will be able to assign the SQL DB Contributor role to other users.                            | <input type="radio"/> | <input type="radio"/> |
| DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server. | <input type="radio"/> | <input type="radio"/> |

Answer:

| Statements  | Yes                              | No                               |
|---|----------------------------------|----------------------------------|
| DBAGroup1 will be able to sign in to each customer's Azure SQL database by using Azure Data Studio.     | <input checked="" type="radio"/> | <input type="radio"/>            |
| DBAGroup1 will be able to assign the SQL DB Contributor role to other users.                            | <input type="radio"/>            | <input checked="" type="radio"/> |
| DBAGroup2 will be able to create a new Azure SQL database on each customer's Azure SQL Database server. | <input checked="" type="radio"/> | <input type="radio"/>            |

Explanation:

Box 1: Yes

DBAGroup1 is member of the Contributor role.

The Contributor role grants full access to manage all resources, but does not allow you to assign roles in Azure RBAC, manage assignments in Azure Blueprints, or share image galleries.

Box 2: No

Box 3: Yes

DBAGroup2 is member of the SQL DB Contributor role.

The SQL DB Contributor role lets you manage SQL databases, but not access to them. Also, you can't manage their security-related policies or their parent SQL servers. As a member of this role you can create and manage SQL databases.

Question: 89

You have 40 Azure SQL databases, each for a different customer. All the databases reside on the same Azure SQL Database server.



You need to ensure that each customer can only connect to and access their respective database.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Implement row-level security (RLS).
- B. Create users in each database.
- C. Configure the database firewall.
- D. Configure the server firewall.
- E. Create logins in the master database.
- F. Implement Always Encrypted.

**Answer:** B,C

Explanation:

Manage database access by adding users to the database, or allowing user access with secure connection strings.

Database-level firewall rules only apply to individual databases.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/secure-database-tutorial>

**Question:** 90

DRAG DROP

You need to implement statistics maintenance for SalesSQLDb1. The solution must meet the technical requirements.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

## Actions

## Answer Area

Create and configure a schedule.

Create a SQL Server Agent job.

Publish the runbook.

Create an Azure Automation account.

Import the SqlServer module.

Create a runbook that runs a PowerShell script.

Run `sp_add_jobserver`.

**Answer:**



## Actions

## Answer Area

|   |   |
|---|---|
| Create and configure a schedule.                | Create an Azure Automation account.             |
| Create a SQL Server Agent job.                  | Import the SqlServer module.                    |
| Publish the runbook.                            | Create a runbook that runs a PowerShell script. |
| Create an Azure Automation account.             | Create and configure a schedule.                |
| Import the SqlServer module.                    |   |
| Create a runbook that runs a PowerShell script. |   |
| Run <code>sp_add_jobserver</code> .             |   |

Explanation:

Automating Azure SQL DB index and statistics maintenance using Azure Automation:

Question: 91

HOTSPOT

You have SQL Server on an Azure virtual machine that contains a database named DB1.

The database reports a CHECKSUM error.

You need to recover the database.

How should you complete the statements? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

USE master;

ALTER DATABASE [DB1] SET

GO

▼

OFFLINE

ONLINE

SINGLE\_USER

TRUSTWORTHY

WITH ROLLBACK IMMEDIATE;

DBCC CHECKDB ('DB1',

GO

▼

MOINDEX

PHYSICAL\_ONLY

REPAIR\_ALLOW\_DATA\_LOSS

REPAIR\_FAST

WITH NO\_INFOMSGS;

ALTER DATABASE [DB1] SET

GO

▼

MULTI\_USER;

ONLINE;

OPEN;

TRUSTWORTHY;

Answer:  
Answer Area

USE master;

ALTER DATABASE [DB1] SET

GO

▼

OFFLINE

ONLINE

SINGLE\_USER

TRUSTWORTHY

WITH ROLLBACK IMMEDIATE;

DBCC CHECKDB ('DB1',

GO

▼

MOINDEX

PHYSICAL\_ONLY

REPAIR\_ALLOW\_DATA\_LOSS

REPAIR\_FAST

WITH NO\_INFOMSGS;

ALTER DATABASE [DB1] SET

GO

▼

MULTI\_USER;

ONLINE;

OPEN;

TRUSTWORTHY;

Explanation:

Box 1: SINGLE\_USER

The specified database must be in single-user mode to use one of the following repair options.

#### Box 2: REPAIR\_ALLOW\_DATA\_LOSS

REPAIR\_ALLOW\_DATA\_LOSS tries to repair all reported errors. These repairs can cause some data loss.

Note: The REPAIR\_ALLOW\_DATA\_LOSS option is a supported feature but it may not always be the best option for bringing a database to a physically consistent state. If successful, the REPAIR\_ALLOW\_DATA\_LOSS option may result in some data loss. In fact, it may result in more data lost than if a user were to restore the database from the last known good backup.

#### Question: 92

You have an instance of SQL Server on Azure Virtual Machine named SQL1.

You need to monitor SQL1 and query the metrics by using Kusto query language. The solution must minimize administrative effort.

Where should you store the metrics?

- A. a Log Analytics workspace
- B. Azure Event Hubs
- C. Azure SQL Database
- D. an Azure Blob storage container

**Answer: A**

#### Question: 93

Topic 5, Misc. Questions

#### HOTSPOT

You have SQL Server on an Azure virtual machine that contains a database named Db1.

You need to enable automatic tuning for Db1.

How should you complete the statements? To answer, select the appropriate answer in the answer area. NOTE: Each correct selection is worth one point.

## Answer Area

ALTER DATABASE [Db1]

|   |   |
|---|---|
|   | ▼ |
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF) |   |
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)  |   |
| SET AUTOMATIC_TUNING=AUTO                       |   |
| SET QUERY_STORE=OFF                             |   |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)    |   |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)   |   |

GO

ALTER DATABASE [Db1]

|   |   |
|---|---|
|   | ▼ |
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF) |   |
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)  |   |
| SET AUTOMATIC_TUNING=AUTO                       |   |
| SET QUERY_STORE=OFF                             |   |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)    |   |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)   |   |

GO

Answer:

## Answer Area

ALTER DATABASE [Db1]

|   |
|---|
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF) |
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)  |
| SET AUTOMATIC_TUNING=AUTO                       |
| SET QUERY_STORE=OFF                             |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)    |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)   |

GO

ALTER DATABASE [Db1]

|   |
|---|
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=OFF) |
| SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN=ON)  |
| SET AUTOMATIC_TUNING=AUTO                       |
| SET QUERY_STORE=OFF                             |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_ONLY)    |
| SET QUERY_STORE=ON(OPERATION_MODE=READ_WRITE)   |

GO

Explanation:

Box 1: SET AUTOMATIC\_TUNING = AUTO

To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:

ALTER DATABASE current SET AUTOMATIC\_TUNING = AUTO

Setting automatic tuning to AUTO will apply Azure Defaults.

Box 2: SET AUTOMATIC\_TUNING (FORCE\_LAST\_GOOD\_PLAN = ON)

To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:

ALTER DATABASE current SET AUTOMATIC\_TUNING (FORCE\_LAST\_GOOD\_PLAN = ON)

Setting the individual tuning option to ON will override any setting that database inherited and enable the tuning option. Setting it to OFF will also override any setting that database inherited and disable the tuning option.

### Question: 94

You are designing a security model for an Azure Synapse Analytics dedicated SQL pool that will support multiple companies.

You need to ensure that users from each company can view only the data of their respective company.

Which two objects should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. a column encryption key
- B. asymmetric keys
- C. a function
- D. a custom role-based access control (RBAC) role
- E. a security policy

**Answer:** D,E

Explanation:

Azure RBAC is used to manage who can create, update, or delete the Synapse workspace and its SQL pools, Apache Spark pools, and Integration runtimes.

Define and implement network security configurations for resources related to your dedicated SQL pool with Azure Policy.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/security/synapse-workspace-synapse-rbac>

<https://docs.microsoft.com/en-us/security/benchmark/azure/baselines/synapse-analytics-security-baseline>

**Question:** 95

DRAG DROP

You have an Azure SQL database named DB1. DB1 contains a table that has a column named Col1.

You need to encrypt the data in Col1.

Which four actions should you perform for DB1 in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



Actions

Answer Area

- Create a database master key.
- Create a column master key.
- Open the symmetric key.
- Create a certificate.
- Update Col1.
- Create a symmetric key.

Answer:

Actions

Answer Area

- |                               |                               |
|-------------------------------|-------------------------------|
| Create a database master key. | Create a database master key. |
| Create a column master key.   | Create a certificate.         |
| Open the symmetric key.       | Create a symmetric key.       |
| Create a certificate.         | Open the symmetric key.       |
| Update Col1.                  |                               |
| Create a symmetric key.       |                               |

Explanation:

Table

Description automatically generated

Question: 96

DRAG DROP



You have SQL Server 2019 on an Azure virtual machine that contains an SSISDB database.

A recent failure causes the master database to be lost.

You discover that all Microsoft SQL Server integration Services (SSIS) packages fail to run on the virtual machine.

Which four actions should you perform in sequence to resolve the issue? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct.

**Actions**

**Answer Area**

Add a certificate to an Azure key vault

Enable Transparent Data Encryption (TDE)

Encrypt a copy of the master key by using the service master key

Turn on the TRUSTWORTHY property and the CLR property

Attach the SSISDB database

Open the master key for the SSISDB database

**Answer:**

## Actions

## Answer Area

Add a certificate to an Azure key vault

Attach the SSISDB database

Enable Transparent Data Encryption (TDE)

Turn on the TRUSTWORTHY property and the CLR property

Encrypt a copy of the master key by using the service master key

Open the master key for the SSISDB database

Turn on the TRUSTWORTHY property and the CLR property

Encrypt a copy of the master key by using the service master key

Attach the SSISDB database

Open the master key for the SSISDB database

Explanation:

Step 1: Attach the SSISDB database

Step 2: Turn on the TRUSTWORTHY property and the CLR property

If you are restoring the SSISDB database to an SQL Server instance where the SSISDB catalog was never created, enable common language runtime (clr)

Step 3: Open the master key for the SSISDB database

Restore the master key by this method if you have the original password that was used to create SSISDB.

open master key decryption by password = 'LS1Setup!' –'Password used when creating SSISDB'

Alter Master Key Add encryption by Service Master Key

Step 4: Encrypt a copy of the mater key by using the service master key

Question: 97

DRAG DROP

You need to implement statistics maintenance for SalesSQLDb1. The solution must meet the technical requirements.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

## Actions

## Answer Area

Create and configure a schedule.

Create a SQL Server Agent job.

Publish the runbook.

Create an Azure Automation account.

Import the SqlServer module.

Create a runbook that runs a PowerShell script.

Run `sp_add_jobserver`.

**Answer:**

## Actions

## Answer Area

|   |   |
|---|---|
| Create and configure a schedule.                | Create an Azure Automation account.             |
| Create a SQL Server Agent job.                  | Import the SqlServer module.                    |
| Publish the runbook.                            | Create a runbook that runs a PowerShell script. |
| Create an Azure Automation account.             | Create and configure a schedule.                |
| Import the SqlServer module.                    |   |
| Create a runbook that runs a PowerShell script. |   |
| Run <code>sp_add_jobserver</code> .             |   |

Explanation:

Automating Azure SQL DB index and statistics maintenance using Azure Automation:

**Question: 98**

You are designing a streaming data solution that will ingest variable volumes of data.

You need to ensure that you can change the partition count after creation.

Which service should you use to ingest the data?

- A. Azure Event Hubs Standard
- B. Azure Stream Analytics
- C. Azure Data Factory
- D. Azure Event Hubs Dedicated

**Answer: D**

Explanation:

The partition count for an event hub in a dedicated Event Hubs cluster can be increased after the event hub has been created.

Reference: <https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

### Question: 99

#### DRAG DROP

You are building an Azure virtual machine.

You allocate two 1-TiB, P30 premium storage disks to the virtual machine. Each disk provides 5,000 IOPS.

You plan to migrate an on-premises instance of Microsoft SQL Server to the virtual machine. The instance has a database that contains a 1.2-TiB data file. The database requires 10,000 IOPS.

You need to configure storage for the virtual machine to support the database.

Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

#### Actions

#### Answer Area

a virtual disk that uses the stripe layout

a virtual disk that uses the mirror layout

a volume

a virtual disk that uses the simple layout

a storage pool

### Answer:

## Actions

## Answer Area

a virtual disk that uses the stripe layout

a storage pool

a virtual disk that uses the mirror layout

a virtual disk that uses the stripe layout

a volume

a volume

a virtual disk that uses the simple layout

a storage pool

Explanation:

Follow these same steps to create striped virtual disk:

Create Log Storage Pool.

Create Virtual Disk

Create Volume

Box 1: a storage pool

Box 2: a virtual disk that uses stripe layout

Disk Striping: Use multiple disks and stripe them together to get a combined higher IOPS and Throughput limit. The combined limit per VM should be higher than the combined limits of attached premium disks.

Box 3: a volume





# SAMPLE QUESTIONS

*These questions are for demo purpose only. **Full version** is up to date and contains actual questions and answers.*

*Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:*

**Actual Exam Questions:** *Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

**Exam Dumps:** *Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

**Practice Tests:** *Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

**Guaranteed Success:** *Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

**Updated Content:** *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

**Technical Support:** *Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*

For More exams visit <https://killexams.com/vendors-exam-list>  
Kill your exam at First Attempt....Guaranteed!