

ITPassLeader



Pass Your Next Certification Exam Fast!

Select a vendor... Select an test... Your email address [Free Download Demo](#)



Instant Download



365 Days Free Updates



Money Back Guarantee



Security & Privacy

Choose the version that fits your needs

PDF Version

Desktop Test Engine

Online Test Engine

Latest and Up-to-Date exam dumps with real exam questions answers.



Get 12-Months free updates without any extra charges.



Experience same exam environment before appearing in the certification exam.



100% exam passing guarantee in the first attempt.



20% discount on more than one license and 30% discount on 5+ license purchases.



100% secure purchase on SSL.



Completely private purchase without sharing your personal info with anyone.



<http://www.itpassleader.com>

High-praise Exam Dumps Questions grant you success by high pass rate - ITPassLeader

Exam : **070-765**

Title : Provisioning SQL Databases

Vendor : Microsoft

Version : DEMO

NO.1 You administer a Microsoft SQL Server 2014 server. The MSSQLSERVER service uses a domain account named CONTOSO\SQLService.

You plan to configure Instant File Initialization.

You need to ensure that Data File Autogrow operations use Instant File Initialization.

What should you do? Choose all that apply.

- A. Restart the SQL Server Agent Service.
- B. Add the CONTOSO\SQLService account to the Server Operators fixed server role.
- C. Enable snapshot isolation.
- D. Add the CONTOSO\SQLService account to the Perform Volume Maintenance Tasks local security policy.
- E. Disable snapshot isolation.
- F. Restart the SQL Server Service.

Answer: D,F

Explanation:

How To Enable Instant File Initialization

References:

<http://msdn.microsoft.com/en-us/library/ms175935.aspx>

NO.2 Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have a virtual machine (VM) in Microsoft Azure, which has a 2 terabyte (TB) database. Microsoft SQL Server backups are performed by using Backup to URL.

You need to provision the storage account for the backups while minimizing costs.

Which storage option should you use?

- A. Premium P30 disk storage
- B. Standard zone redundant blob storage
- C. Standard locally redundant blob storage
- D. Standard locally redundant disk storage
- E. Standard geo-redundant disk storage
- F. Premium P10 disk storage
- G. Standard geo-redundant blob storage
- H. Premium P20 disk storage

Answer: C

Explanation:

A URL specifies a Uniform Resource Identifier (URI) to a unique backup file. The URL is used to provide the location and name of the SQL Server backup file. The URL must point to an actual blob, not just a container. If the blob does not exist, it is created. If an existing blob is specified, BACKUP fails, unless the "WITH FORMAT" option is specified to overwrite the existing backup file in the blob. LOCALLY REDUNDANT STORAGE (LRS) makes multiple synchronous copies of your data within a single datacenter.

NO.3 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 use a Microsoft Azure SQL database as a data warehouse. The database is in the Standard service tier and has 400 elastic database throughput units (eDTUs).

You load data to the database by using Azure Data Factory.

You need to reduce the amount of time it takes to load the data.

Solution: You move the database to a Standard elastic pool that has 800 eDTUs.

Does this meet the goal?

A. Yes

B. No

Answer: B

NO.4 Background

You manage the Microsoft SQL Server environment for a company that manufactures and sells automobile parts.

The environment includes the following servers: SRV1 and SRV2. SRV1 has 16 logical cores and hosts a SQL Server instance that supports a mission-critical application. The application has approximately 30,000 concurrent users and relies heavily on the use of temporary tables.

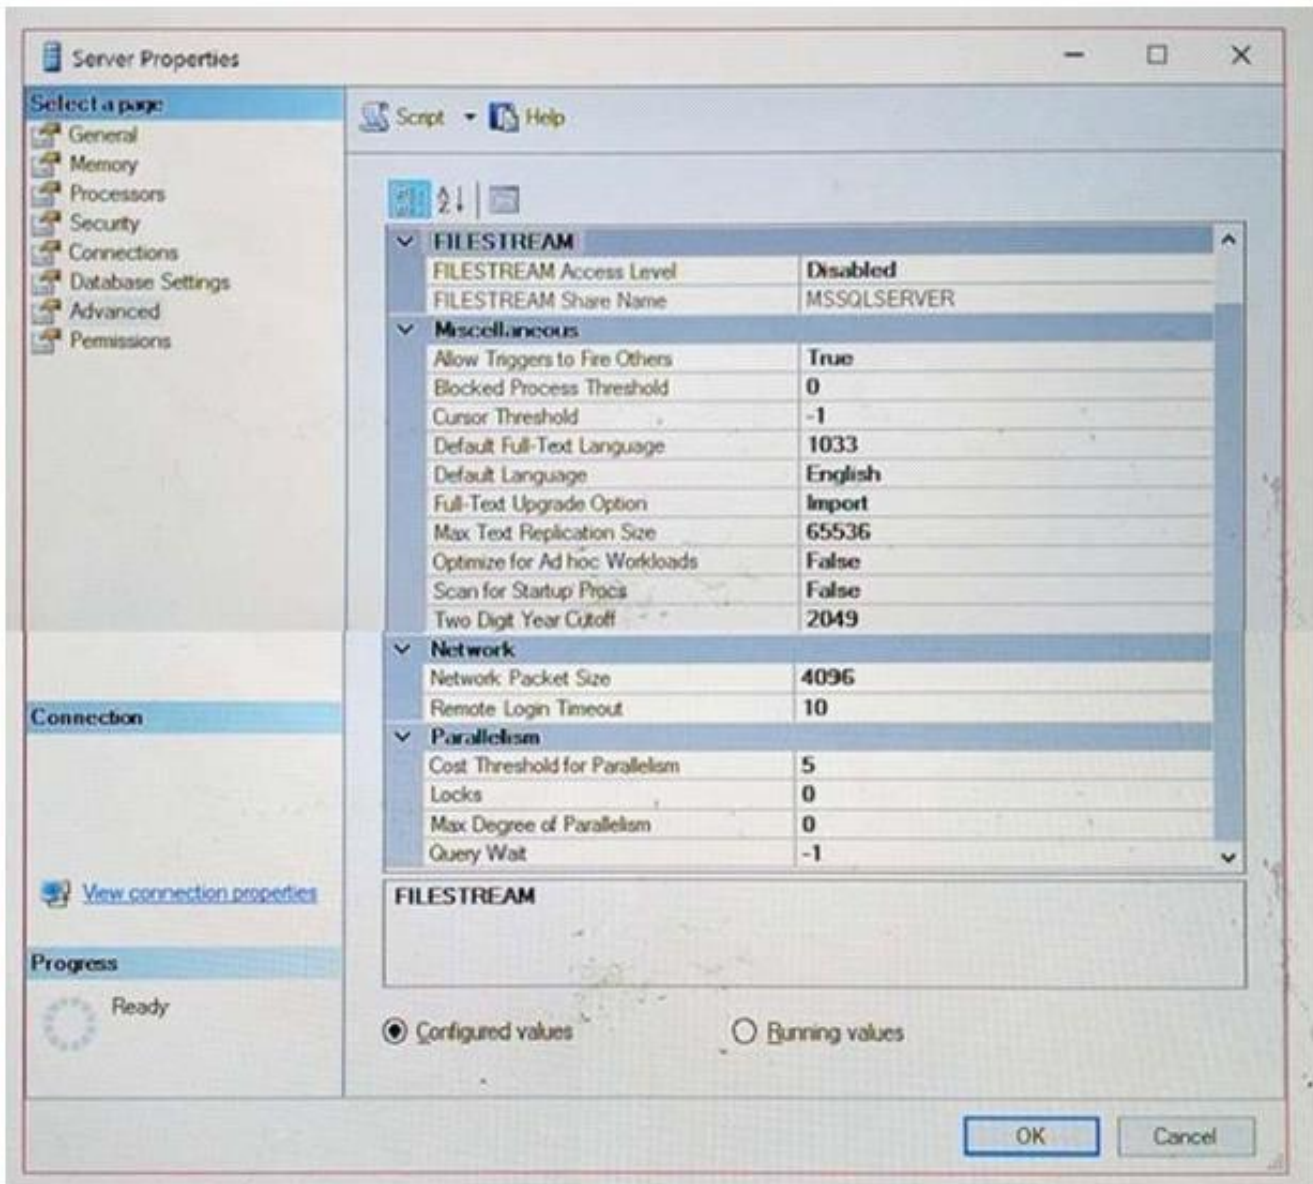
The environment also includes the following databases: DB1, DB2, and Reporting. The Reporting database is protected with Transparent Data Encryption (TDE). You plan to migrate this database to a new server. You detach the database and copy it to the new server.

You are performing tuning on a SQL Server database instance. The application which uses the database was written using an object relationship mapping (ORM) tool which maps tables as objects within the application code. There are 30 stored procedures that are regularly used by the application.

After reviewing the plan cache you have identified that a large number of simple queries are using parallelism, and that execution plans are not being kept in the plan cache for very long.

You review the properties of the instance (Click the Exhibit button).

Exhibit:



You need to restore the Reporting database to SRV2. What should you do? To answer, drag the appropriate options to the correct locations. Each option 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.

Select and Place:

Values

master encryption key on the master database

service master key

server certificate

Reporting database .mdf file

master key password

Answer area

1. Copy the certificate and private key backups from the old server to the new server.

2. Create:

3. Restore:

4. Attach the Reporting database.

Answer:

Values
 master encryption key on the master database

 service master key

 server certificate

 Reporting database .mdf file

 master key password
Answer area

1. Copy the certificate and private key backups from the old server to the new server.

2. Create: server certificate

3. Restore: Reporting database .mdf file

4. Attach the Reporting database.

Explanation:

Step 2: Create: server certificate

Recreate the server certificate by using the original server certificate backup file.

Note: The password must be the same as the password that was used when the backup was created.

Step 3: Restore: Reporting database .mdf file.

-- Attach the database that is being moved.

-- The path of the database files must be the location where you have stored the database files.

Example:

```
CREATE DATABASE [CustRecords] ON
```

```
( FILENAME = N'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\CustRecords.mdf' ),
```

```
( FILENAME = N'C:\Program Files\Microsoft SQL Server\MSSQL13.MSSQLSERVER\MSSQL\DATA\CustRecords_log.LDF' )
```

```
FOR ATTACH ;
```

```
GO
```

From scenario: The Reporting database is protected with Transparent Data Encryption (TDE). You plan to migrate this database to a new server. You detach the database and copy it to the new server

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/move-a-tdeprotected-database-to-another-sql-server>

NO.5 You have an on-premises database.

You plan to migrate the database to Microsoft SQL Server on a Microsoft Azure virtual machine.

You move the database files to Azure.

You need to attach the database files to the SQL Server instance on the virtual machine. The solution must ensure that you can run file snapshot backups.

How should you complete the statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer area

```
USE (master)
GO
CREATE DATABASE [Production_DB]
(
    [ ] = N'https://proddbstorage=contoso.blob.core.windows.net/datafiles/proddb.mdf'
    [ ]
    [ ]
    [ ]
    [ ]
)
(
    [ ]
    [ ]
)
GO
CREATE
```

Answer:

Answer area

```
USE (master)
GO
CREATE DATABASE [Production_DB]
(
    [ ] = N'https://proddbstorage=contoso.blob.core.windows.net/datafiles/proddb.mdf'
    [ ]
    [ ]
    [ ]
)
(
    [ ]
    [ ]
)
GO
CREATE
```

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-sql-server-transact-sql>

NO.6 Your company has several Microsoft Azure SQL Database instances used within an elastic pool. You need to obtain a list of databases in the pool.

How should you complete the commands? To answer, drag the appropriate segments to the correct targets. Each segment 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.

Segments

- elastic-pool
- list
- list-dbs
- list-editions

Answer Area

az sql [Segment] [Segment]

Answer:

Segments

elastic-pool
list
list-dbs
list-editions

Answer Area

az sql

elastic-pool

list-dbs

Explanation:

References:

<https://docs.microsoft.com/en-us/cli/azure/sql/elastic-pool?view=azure-cli-latest#az-sql-elastic-pool-list-dbs>

NO.7 You plan to migrate a Microsoft sql server instance between physical servers.

You must migrate the metadata associated with the database instance.

You need to ensure that the new instance retains the existing jobs and alerts.

Solutions: You restore the model database.

Does the solution meet the goal?

A. No

B. Yes

Answer: A

Explanation:

The model database does not handle alerts and jobs. It is used as the template for all databases created on an instance of SQL Server.

The msdb database is used by SQL Server Agent for scheduling alerts and jobs and by other features such as SQL Server Management Studio, Service Broker and Database Mail.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/databases/msdb-database?view=sql-server-2017>

NO.8 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution. Determine whether the solution meets stated goals.

You manage a Microsoft SQL Server environment with several databases.

You need to ensure that queries use statistical data and do not initialize values for local variables.

Solution: You enable the LEGACY_CARDINALITY_ESTIMATION option for the databases.

Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }

Enables you to set the query optimizer cardinality estimation model to the SQL Server 2012 and

earlier version independent of the compatibility level of the database. This is equivalent to Trace Flag 9481.

References:

<https://msdn.microsoft.com/en-us/library/mt629158.aspx>

NO.9 A company has an on-premises Microsoft SQL Server 2014 environment. The company has a main office in Seattle, and remote offices in Amsterdam and Tokyo. You plan to deploy a Microsoft Azure SQL Database instance to support a new application. You expect to have 100 users from each office.

In the past, users at remote sites reported issues when they used applications hosted at the Seattle office.

You need to optimize performance for users running reports while minimizing costs.

What should you do?

- A.** Deploy a database from the Premium service tier.
- B.** Implement replication from an on-premises SQL Server database to the Azure SQL Database instance.
- C.** Implement a standard database with readable secondaries in Asia and Europe, and then migrate the application.
- D.** Implement an elastic pool.

Answer: C

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-geo-replication-transact-sql#add-secondary-database>

NO.10 Note: This questions is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You deploy Microsoft SQL Server to a virtual machine in Azure. You distribute the database files and filegroups across multiple Azure storage disks.

You must be able to manage the databases as individual entities by using SQL Server Management Studio. All data in the databases must be stored encrypted. Backups must be encrypted by using the same key as the live copy of the database.

You need to secure the data.

What should you implement?

- A.** BitLocker
- B.** transport-level encryption
- C.** Transparent Data Encryption
- D.** dynamic data masking
- E.** Always Encrypted
- F.** cell-level encryption
- G.** Encrypting File System

Answer: C

Explanation:

Transparent data encryption (TDE) encrypts your databases, associated backups, and transaction log files at rest without requiring changes to your applications.

TDE encrypts the storage of an entire database by using a symmetric key called the database encryption key. In SQL Database the database encryption key is protected by a built-in server certificate. The built-in server certificate is unique for each SQL Database server.

References:

<https://msdn.microsoft.com/en-us/library/dn948096.aspx>

NO.11 You have Microsoft SQL Server on a DS-series Microsoft Azure virtual machine. The virtual machine has a production database named DB1. All database files are on drive E and use standard storage.

Users report that queries take a long time to execute.

You discover that the queries are waiting for pagelatch_IO.

You need to reduce the amount of time it takes for the queries to execute.

What should you do?

- A. Move the msdb databases to drive D.
- B. Move the databases for DB1 to drive D.
- C. On drive E add more database files for DB1.
- D. Change drive E to Premium Storage.

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disks-types>

Answer: D

NO.12 You plan to deploy Microsoft SQL Server on a Microsoft Azure Virtual machine. The virtual machine will have a 30-TB database and will have 10 1-TB VHDs for the database.

You need to configure the storage to meet the following requirements:

Evenly distribute read and write operations across the VHDs.

Minimize the read and write time.

Which storage configuration should you use?

- A. a striped volume
- B. a RAID-5 volume
- C. a simple storage pool
- D. a mirrored storage pool
- E. a parity storage pool

Answer: A

Explanation:

Data that is written to a striped volume is interleaved to all disks at the same time instead of sequentially.

Therefore, disk performance is the fastest on a RAID 0 volume as compared to any other type of disk configuration.

NO.13 You upgrade a database named DB1 to Microsoft SQL Server 2016.

Users report that performance for several queries is degraded. You determine that the query optimizer is incorrectly estimating the number of rows that the queries will return.

You need to resolve the performance issues.

How should you complete the Transact-SQL statement? To answer, drag the Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

Transact-SQL segments

SCOPED CONFIGURATION
 LEGACY_CARDINALITY_ESTIMATION = ON
 LEGACY_CARDINALITY_ESTIMATION = OFF
 SCOPED CONFIGURATION FOR SECONDARY
 QUERY_STORE = ON
 QUERY_STORE = OFF

Answer area

```
ALTER DATABASE DB1
  Transact-SQL segment
SET
  Transact-SQL segment
```

Answer:**Transact-SQL segments**

SCOPED CONFIGURATION
 LEGACY_CARDINALITY_ESTIMATION = ON
 LEGACY_CARDINALITY_ESTIMATION = OFF
 SCOPED CONFIGURATION FOR SECONDARY
 QUERY_STORE = ON
 QUERY_STORE = OFF

Answer area

```
ALTER DATABASE DB1
  SCOPED CONFIGURATION
SET
  LEGACY_CARDINALITY_ESTIMATION = ON
```

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql?view=sql-server-2017>

NO.14 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are tuning the performance of a virtual machines that hosts a Microsoft SQL Server instance.

The virtual machine originally had four CPU cores and now has 32 CPU cores.

The SQL Server instance uses the default settings and has an OLTP database named db1. The largest table in db1 is a key value store table named table1.

Several reports use the PIVOT statement and access more than 100 million rows in table1.

You discover that when the reports run, there are PAGELATCH_IO waits on PFS pages 2:1:1, 2:2:1, 2:3:1, and 2:4:1 within the tempdb database.

You need to prevent the PAGELATCH_IO waits from occurring.

Solution: You add more files to db1.

Does this meet the goal?

A. Yes

B. No

Answer: A

Explanation:

From SQL Server's perspective, you can measure the I/O latency from sys.dm_os_wait_stats. If you consistently see high waiting for PAGELATCH_IO, you can benefit from a faster I/O subsystem for SQL Server.

A cause can be poor design of your database - you may wish to split out data located on 'hot pages', which are accessed frequently and which you might identify as the causes of your latch contention.

For example, if you have a currency table with a data page containing 100 rows, of which 1 is updated per transaction and you have a transaction rate of 200/sec, you could see page latch queues of 100 or more. If each page latch wait costs just 5ms before clearing, this represents a full half-second delay for each update. In this case, splitting out the currency rows into different tables might prove more performant (if less normalized and logically structured).

References:

<https://www.mssqltips.com/sqlservertip/3088/explanation-of-sql-server-io-and-latches/>

NO.15 Your company has many Microsoft SQL Server instances hosted in a data center. You also manage five Microsoft Azure SQL Database instances that are hosted on a single server in Azure. You need to minimize costs associated with Azure resources while maintaining the current performance levels of each Azure SQL Database instance.

Which three 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

Determine the maximum DTU usage of all Azure SQL databases.

Add databases to the pool.

Add a new pool to the Azure SQL server.

Create an Azure Availability group.

Create an Azure Blob store.

Create an Azure SQL server.

Answer Area

Answer:

Actions

Determine the maximum DTU usage of all Azure SQL databases.

Add databases to the pool.

Add a new pool to the Azure SQL server.

Create an Azure Availability group.

Create an Azure Blob store.

Create an Azure SQL server.

Answer Area

Create an Azure SQL server.

Add a new pool to the Azure SQL server.

Add databases to the pool.

Explanation:

SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single Azure SQL Database server and share a set number of resources at a set price.

References:

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

NO.16 You plan to migrate on-premises Microsoft SQL Server databases to SQL Server on a Microsoft Azure virtual machine.

A full backup of the databases is 3 TB. The corporate network has a 1.5 MB/sec Internet connection. You need to move the backups to Azure to ensure that you can restore the databases to the Azure virtual machine as quickly as possible.

What should you do?

- A. Run azcopy.exe and specify the /SetContentType parameter.
- B. Create an availability group that is configured for automatic seeding.
- C. Use the Azure Import/Export service.
- D. Use SQL Server Backup to URL
- E. Run robocopy.exe and specify the /mir parameter.

Answer: C

Explanation:

Azure Import/Export service is used to securely import large amounts of data to Azure Blob storage and Azure Files by shipping disk drives to an Azure datacenter. This service can also be used to transfer data from Azure Blob storage to disk drives and ship to your on-premises sites. Data from one or more disk drives can be imported either to Azure Blob storage or Azure Files.

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-import-export-service>