

## 70-765 Dumps

### Provisioning SQL Databases (beta)

<https://www.certleader.com/70-765-dumps.html>



**NEW QUESTION 1**

- (Topic 1)

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.

Your company plans to use Microsoft Azure Resource Manager templates for all future deployments of SQL Server on Azure virtual machines.

You need to create the templates.

Solution: You use Visual Studio to create a JSON template that defines the deployment and configuration settings for the SQL Server environment.

Does the solution meet the goal?

A. Yes

B. No

**Answer: A**

**Explanation:**

Azure Resource Manager template consists of JSON, not XAML, and expressions that you can use to construct values for your deployment.

A good JSON editor can simplify the task of creating templates.

Note: In its simplest structure, an Azure Resource Manager template contains the following elements:

```
{
"$schema": "http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
"contentVersion": "", "parameters": { },
"variables": { },
"resources": [ ],
"outputs": { }
}
```

References:<https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-authoring-templates>

**NEW QUESTION 2**

- (Topic 2)

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 manage on-premises and Microsoft Azure SQL Database instances for a company. Your environment must support the Microsoft SQL Server 2012 ODBS driver.

You need to encrypt only specific columns in the database. What should you implement?

A. transport-level encryption

B. cell-level encryption

C. Transparent Data Encryption

D. Always Encrypted

E. Encrypting File System

F. BitLocker

G. dynamic data masking

**Answer: D**

**Explanation:**

To encrypt columns you can configure Always Encrypted.

SQL Server Management Studio (SSMS) provides a wizard that helps you easily configure Always Encrypted by setting up the column master key, column encryption key, and encrypted columns for you.

Always Encrypted allows client applications to encrypt sensitive data and never reveal the data or the encryption keys to SQL Server or Azure SQL Database. An Always Encrypted enabled driver, such as the ODBC Driver 13.1 for SQL Server, achieves this by transparently encrypting and decrypting sensitive data in the client application.

Note: The ODBC driver automatically determines which query parameters correspond to sensitive database columns (protected using Always Encrypted), and encrypts the values of those parameters before passing the data to SQL Server or Azure SQL Database. Similarly, the driver transparently decrypts data retrieved from encrypted database columns in query results.

References:<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault#encrypt-columns-configure-always-encrypted>  
[https://msdn.microsoft.com/en-us/library/mt637351\(v=sql.110\).aspx](https://msdn.microsoft.com/en-us/library/mt637351(v=sql.110).aspx)

**NEW QUESTION 3**

DRAG DROP - (Topic 2)

A new Azure Active Directory security principal named ReportUser@contoso.onmicrosoft.com should have access to select all current and future objects in the Reporting database. You should not grant the principal any other

permissions. You should use your Active Directory Domain Services (AD DS) account to authenticate to the Azure SQL database.

You need to create the new security principal.

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  | Answer Area |
|--|-------------|
| Create a connection to the <b>master</b> database on the Azure SQL Server instance by using your Active Directory authenticated account.   |             |
| Create a connection to the <b>Reporting</b> database on the Azure SQL Server instance by using your Active Directory authenticated account.  |             |
| Run the following Transact-SQL statement:<br><br>EXEC sp_addrolemember 'db_datareader',<br>'reportuser@contoso.onmicrosoft.com'  |             |
| Run the following Transact-SQL statement:<br><br>CREATE USER<br>[reportuser@contoso.onmicrosoft.com]<br>FROM EXTERNAL PROVIDER   |             |
| Run the following Transact-SQL statements:<br><br>USE Reporting<br>CREATE USER<br>[reportuser@contoso.onmicrosoft.com] FOR<br>LOGIN<br>[reportuser@contoso.onmicrosoft.com]<br>GRANT SELECT TO<br>[reportuser@contoso.onmicrosoft.com] |             |
| Create a connection to the <b>Reporting</b> database on the Azure SQL Server instance by using your SQL Server authenticated account.  |             |

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Step 1:

To provision an Azure AD-based contained database user (other than the server administrator that owns the database), connect to the database (here the Reporting database) with an Azure AD identity (not with a SQL Server account) that has access to the database.

Step 2: CREATE USER ... FROM EXTERNAL PROVIDER

To create an Azure AD-based contained database user (other than the server administrator that owns the database), connect to the database with an Azure AD identity, as a user with at least the ALTER ANY USER permission. Then use the following Transact-SQL syntax:

CREATE USER <Azure\_AD\_principal\_name> FROM EXTERNAL PROVIDER;

Step 3:

Grant the proper reading permissions.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-aad-authentication>

**NEW QUESTION 4**

- (Topic 2)

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

**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>

NEW QUESTION 5

- (Topic 3)

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

Answer: B

Explanation:

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

NEW QUESTION 6

HOTSPOT - (Topic 4)

You need to resolve the identified issues.

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

Answer Area

What setting would you change to reduce the number of execution plans in the plan cache?

Optimize for Ad Hoc workload  
Max Degree of Parallelism  
Query Wait

What setting would you change to which value to reduce the number of queries which are using parallelism?

Max Degree of Parallelism to 4  
Cost Threshold for Parallelism to 50  
Locks to 100

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

From exhibit we see:

Cost Threshold of Parallelism: 5 Optimize for Ad Hoc Workloads: false

Max Degree of Parallelism: 0 (This is the default setting, which enables the server to determine the maximum degree of parallelism. It is fine.)

Locks: 0

Query Wait: -1

Box 1: Optimize for Ad Hoc Workload

Change the Optimize for Ad Hoc Workload setting from false to 1/True.

The optimize for ad hoc workloads option is used to improve the efficiency of the plan cache for workloads that contain many single use ad hoc batches. When this option is set to 1, the Database Engine stores a small compiled plan stub in the plan cache when a batch is compiled for the first time, instead of the full compiled plan. This helps to relieve memory pressure by not allowing the plan cache to become filled with compiled plans that are not reused.

NEW QUESTION 7

HOTSPOT - (Topic 4)

You need to optimize SRV1.

What configuration changes should you implement? To answer, select the appropriate option from each list in the answer area.

Answer Area

How should you modify the tempdb configuration?

Change the recovery model of tempdb.  
Change the number of tempdb files.  
Change the size of the tempdb log file.  
Change the MAXDOP property.

How should you reconfigure the tempdb database?

Add additional tempdb files.  
Remove tempdb files.  
Add tempdb log files.  
Remove tempdb log files.  
Set MAXDOP to 8.



- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

From the scenario: 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.

Box 1: Change the size of the tempdb log file.

The size and physical placement of the tempdb database can affect the performance of a system. For example, if the size that is defined for tempdb is too small, part of the system- processing load may be taken up with autogrowing tempdb to the size required to support the workload every time you restart the instance of SQL Server. You can avoid this overhead by increasing the sizes of the tempdb data and log file.

Box 2: Add additional tempdb files.

Create as many files as needed to maximize disk bandwidth. Using multiple files reduces tempdb storage contention and yields significantly better scalability.

However, do not create too many files because this can reduce performance and increase management overhead. As a general guideline, create one data file for each CPU on the server (accounting for any affinity mask settings) and then adjust the number of files up or down as necessary.

**NEW QUESTION 8**

HOTSPOT - (Topic 5)

You need to create the contosodb1 database.

How should you complete the Azure PowerShell command? To answer, select the appropriate Azure PowerShell segments in the answer area.

**Answer Area**

|                        |   |
|------------------------|---|
|                        | ▼ |
| New-AzureSqlDatabase   |   |
| New-AzureRmSqlDatabase |   |
| Set-AzureRmSqlDatabase |   |

- ResourceGroupName “contosodbrg”

- ServerName “contososrv”

-DatabaseName “contosodb1”

- Edition

|          |   |
|----------|---|
|          | ▼ |
| Basic    |   |
| Standard |   |
| Premium  |   |

-RequestedServiceObjectName S2

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: New-AzureRmSqlDatabase

New-AzureRmSqlDatabase creates a database or an elastic database.

New-AzureRmSqlDatabase is a command with the Azure Resource Manager (AzureRM) module. Azure Resource Manager enables you to work with the resources in your solution as a group.

**NEW QUESTION 9**

- (Exam Topic 7)


You have Microsoft SQL Server on a Microsoft Azure virtual machine. The virtual machine has a database named DB1. DB1 contains a table named Table1 that has 4 billion rows.

Users report that a query using Table1 takes longer than expected to execute.

You review the execution plan for the query and discover that the expected number of returned rows is one, while the actual number of returned rows is 1 million.

You need to reduce the amount of time it takes for the query to execute. The solution must prevent additional performance issues from being introduced.

Hot Area:



- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

When you set the AUTO\_CREATE\_STATISTICS option on, the Query Optimizer creates statistics on individual columns used in a predicate, if these statistics are not already available. These statistics are necessary to generate the query plan.

References:

<https://www.mssqltips.com/sqlservertip/2766/sql-server-auto-update-and-auto-create-statisticsoptions/>

**NEW QUESTION 10**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 failover cluster that contains two nodes named Node A and Node B. A single instance of SQL Server is installed on the cluster.

An additional node named Node C has been added to the existing cluster.

You need to ensure that the SQL Server instance can use all nodes of the cluster. What should you do?

- A. Run the New SQL Server stand-alone installation Wizard on Node C.
- B. Run the Add Node to SQL Server Failover Cluster Wizard on Node C.
- C. Use Node B to install SQL Server on Node C.
- D. Use Node A to install SQL Server on Node C.

**Answer: B**

**Explanation:**

To add a node to an existing SQL Server failover cluster, you must run SQL Server Setup on the node that is to be added to the SQL Server failover cluster instance. Do not run Setup on the active node.

The Installation Wizard will launch the SQL Server Installation Center. To add a node to an existing failover cluster instance, click Installation in the left-hand pane. Then, select Add node to a SQL Server failover cluster.

References:

<http://technet.microsoft.com/en-us/library/ms191545.aspx>

**NEW QUESTION 10**

- (Exam Topic 7)

You deploy a new Microsoft Azure SQL database instance to support a variety of mobile application and public websites. You configure geo-replication with regions in Brazil and Japan.

You need to implement real-time encryption of the database and all backups. Solution: you enable Dynamic Data Masking on the primary replica. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

SQL Database dynamic data masking does not encrypt the data. Transparent Data Encryption (TDE) would provide a solution.

Note: SQL Database dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.

References:

<https://azure.microsoft.com/en-us/blog/how-to-configure-azure-sql-database-geo-dr-with-azure-key-vault/>

**NEW QUESTION 14**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database.

The database contains a Product table created by using the following definition:

```
CREATE TABLE dbo.Product
(
    ProductID INT PRIMARY KEY,
    Name VARCHAR(50) NOT NULL,
    Color VARCHAR(15) NOT NULL,
    Size VARCHAR(5) NOT NULL,
    Style CHAR(2) NULL,
    Weight DECIMAL(8,2) NULL);
```

You need to ensure that the minimum amount of disk space is used to store the data in the Product table. What should you do?

- A. Convert all indexes to Column Store indexes.
- B. Implement Unicode Compression.
- C. Implement row-level compression.
- D. Implement page-level compression.

**Answer:** D

**NEW QUESTION 17**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 Enterprise Edition server that uses 64 cores.

You discover performance issues when large amounts of data are written to tables under heavy system load. You need to limit the number of cores that handle I/O.

What should you configure?

- A. Processor affinity
- B. Lightweight pooling
- C. Max worker threads
- D. I/O affinity

**Answer:** D

**Explanation:**

The affinity Input-Output (I/O) mask Server Configuration Option.

To carry out multitasking, Microsoft Windows 2000 and Windows Server 2003 sometimes move process threads among different processors. Although efficient from an operating system point of view, this activity can reduce Microsoft SQL Server performance under heavy system loads, as each processor cache is repeatedly reloaded with data. Assigning processors to specific threads can improve performance under these conditions by eliminating processor reloads; such an association between a thread and a processor is called processor affinity.

References:

<http://msdn.microsoft.com/en-us/library/ms189629.aspx>


**NEW QUESTION 20**

- (Exam Topic 7)

You use a Microsoft SQL Server 2014 database that contains two tables named SalesOrderHeader and SalesOrderDetail. The indexes on the tables are as shown in the exhibit.

(Click the Exhibit button.)





You write the following Transact-SQL query:

```
SELECT h.SalesOrderID, h.TotalDue, d.OrderQty
FROM Sales.SalesOrderHeader AS h
    INNER JOIN Sales.SalesOrderDetail AS d
    ON h.SalesOrderID = d.SalesOrderID
WHERE h.TotalDue > 100
AND (d.OrderQty > 5 OR d.LineTotal < 1000.00);
```

You discover that the performance of the query is slow. Analysis of the query plan shows table scans where the estimated rows do not match the actual rows for SalesOrderHeader by using an unexpected index on SalesOrderDetail.

You need to improve the performance of the query. What should you do?

- A. Use a FORCESCAN hint in the query.
- B. Add a clustered index on SalesOrderID in SalesOrderHeader.
- C. Use a FORCESEEK hint in the query.
- D. Update statistics on SalesOrderID on both tables.

**Answer:** D

**Explanation:**

New statistics would be useful.

The UPDATE STATISTICS command updates query optimization statistics on a table or indexed view. By default, the query optimizer already updates statistics as necessary to improve the query plan; in some cases you can improve query performance by using UPDATE STATISTICS or the stored procedure sp\_updatestats to update statistics more frequently than the default updates.

References:

<http://msdn.microsoft.com/en-us/library/ms187348.aspx>

**NEW QUESTION 24**

- (Exam Topic 7)

You administer a Windows 2008 server hosting an instance of Microsoft SQL Server 2014 Standard Edition. The server hosts a database named Orders.

Users report that a query that filters on OrderDate is taking an exceptionally long time. You discover that an index named IX\_OrderDate on the CustomerOrder table is heavily fragmented.

You need to improve the performance of the IX\_OrderDate index. The index should remain online during the operation.

Which Transact-SQL command should you use?



- A. ALTER INDEX IX\_OrderDateON CustomerOrder DISABLE
- B. ALTER INDEX IX\_OrderDateON CustomerOrder ENABLE
- C. ALTER INDEX IX\_OrderDateON CustomerOrder REORGANIZE
- D. ALTER INDEX IX\_OrderDateON CustomerOrder REBUILD

**Answer: C**

**Explanation:**

Reorganize: This option is more lightweight compared to rebuild. It runs through the leaf level of the index, and as it goes it fixes physical ordering of pages and also compacts pages to apply any previously set fillfactor settings. This operation is always online, and if you cancel it then it's able to just stop where it is (it doesn't have a giant operation to rollback).

References:<https://www.brentozar.com/archive/2013/09/index-maintenance-sql-server-rebuild-reorganize/>

**NEW QUESTION 26**

- (Exam Topic 7)

You have been hired as a Database Consultant by ABC.com to design a SQL Server 2014 database solution. You are tasked with designing a scale-out and high-availability SQL Server 2014 Online Transaction

Processing (OLTP) database solution that will maintain copies of data across two server instances.

Your solution must provide scale-out of read operations by distributing the reads from clients across two SQL Server 2014 nodes. The data in both SQL Server nodes needs to be indexed.

What should you include in your solution?

- A. You should include a primary database with scheduled log shipping to the secondary database configured.
- B. You should include two servers configured in an Active-Passive SQL Server 2014 Cluster.
- C. You should include a primary SQL Server 2014 database that uses transactional replication to replicate data to the secondary database.
- D. You should include two servers in an Asynchronous-Commit Availability Mode Availability Group.
- E. You should include two servers in a Synchronous-Commit Availability Mode Availability Group.

**Answer: C**

**Explanation:**

Peer-to-peer replication provides a scale-out and high-availability solution by maintaining copies of data across multiple server instances, also referred to as nodes.

Built on the foundation of transactional replication, peer-to-peer replication propagates transactionally consistent changes in near real-time. This enables applications that require scale-out of read operations to distribute the reads from clients across multiple nodes. Because data is maintained across the nodes in near real-time, peer-to-peer replication provides data redundancy, which increases the availability of data.

References:[https://msdn.microsoft.com/en-us/library/ms151196\(v=sql.110\)](https://msdn.microsoft.com/en-us/library/ms151196(v=sql.110))

**NEW QUESTION 27**

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account.

You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

Solution: You run the Get-Counter cmdlet and specify the -counter '\physicaldisk:disk Transfers/sec' parameter.

Does this meet the goal?

- A. Yes
- B. No

**Answer: A**

**NEW QUESTION 29**

- (Exam Topic 7)

You have an on-premises database that runs several maintenance jobs. You move the database to a Microsoft Azure SQL database.

You need to ensure that the maintenance jobs on indexes continue to run after the move.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

**Actions, Select from these**

**Answer Area, Place here**

Create a runbook

1.

Create an Automation Account

2.

Configure a schedule

3.

Create a credential

4.

Publish a runbook

5.

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

General steps for a solution to automate the maintenance you Azure SQL DB statistics: References:

<https://blogs.msdn.microsoft.com/azuresqlsupport/2018/01/15/automating-azure-sql-db-index-and-statistic-m>

**NEW QUESTION 30**

- (Exam Topic 7)

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'
    )
    (
        ON PRIMARY;
        ON COLLATE;
    )
)
GO
CREATE
```

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 33**

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account.

You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

Solution: You run the Get-Counter cmdlet and specify the -counter '\physicaldisk:disk write/sec' parameter. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**NEW QUESTION 38**

- (Exam Topic 7)

You have a database named DB1 that contains a table named Table1. Table1 has 1 billion rows.

You import 10 million rows of data into Table1. After the import, users report that queries take longer than usual to execute.

You need to identify whether an out-of-date execution plan is causing the performance issue. Which dynamic management view should you use?

- A. sys.dm\_xtp\_transaction\_stats
- B. sys.dm\_exec\_input\_buffer
- C. sys.dm\_db\_index\_operational\_stats
- D. sys.dm\_db\_stats\_properties

**Answer:** C

**Explanation:**

sys.dm\_db\_index\_operational\_stats dynamic management function provides us the current low-level I/O, locking, latching, and access method for each partition of the table. This information is really useful to troubleshoot SQL Server performance issues.

Reference:

<https://basitaalishan.com/2013/03/19/using-sys-dm-db-index-operational-stats-to-analyse-howindexes-are-utili>

**NEW QUESTION 43**

- (Exam Topic 7)

You are the database administrator in your company. You plan to create 10 identical environments that use SQL Server 2016 as a database engine. Each environment has the following custom requirements:

Three user databases must be preinstalled.

The tempdb database must contain eight data files that are 1024 MB each.

Trace flag 2371 must be turned at the instance level.

The solution must meet the following requirements:

The instance must be preconfigured.

No other database features are required in the future.

The solution must use the minimum administrative effort.

You need to prepare the environments. What should you do?

- A. Provision 10 Azure virtual machines that each contain SQL Server 2016, installed by using the default settings.
- B. Create an installation configuration file and perform unattended installations of SQL Server 2016.
- C. Create a virtual machine template by using a prepared instance of SQL Server 2016.
- D. Create a virtual machine template by using a complete instance of SQL Server 2016.

**Answer:** D

**Explanation:**

You should create a virtual machine template by using a complete instance of SQL Server 2016. You use the sysprep tool to prepare a complete instance of SQL Server 2016. By using a complete instance, SQL Server, the network, and the users are all created, and the system cannot be reconfigured during the installation process.

**NEW QUESTION 44**

- (Exam Topic 7)

You administer a SQL Server 2014 server that contains a database named SalesDb. SalesDb contains a schema named Customers that has a table named Regions. A user named UserA is a member of a role named Sales. UserA is granted the Select permission on the Regions table. The Sales role is granted the Select permission on the Customers schema.

You need to ensure that the following requirements are met: Which Transact-SQL statement should you use?

- A. REVOKE SELECT ON Schema::Customers FROM UserA
- B. DENY SELECT ON Object::Regions FROM UserA
- C. EXEC sp\_addrolemember 'Sales', 'UserA'
- D. DENY SELECT ON Object::Regions FROM Sales
- E. REVOKE SELECT ON Object::Regions FROM UserA
- F. DENY SELECT ON Schema::Customers FROM Sales
- G. DENY SELECT ON Schema::Customers FROM UserA
- H. EXEC sp\_droprolemember 'Sales', 'UserA'
- I. REVOKE SELECT ON Object::Regions FROM Sales
- J. REVOKE SELECT ON Schema::Customers FROM Sales

**Answer:** J

**Explanation:**

Use REVOKE to remove the grant or deny of a permission.

References:<https://docs.microsoft.com/en-us/sql/t-sql/statements/permissions-grant-deny-revoke-azure-sql-data->

**NEW QUESTION 45**

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft azure virtual machine that has 12 databases. All database files are in the same Azure Blob storage account.

You need to receive an email notification if I/O operations to the database files exceed 800 MB/s for more than five minutes.

Solution: You run the Add-AzureRmMetricAlertRule cmdlet and specify the –MetricName ‘Network Out’ parameter.

Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**NEW QUESTION 47**

- (Exam Topic 7)

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

**Answer:** D

**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.

Reference:

<https://support.microsoft.com/en-us/help/323433/how-to-establish-a-striped-volume-raid-0-inwindows-server-20>

**NEW QUESTION 49**

- (Exam Topic 7)

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 migrating an on-premises Microsoft SQL Server instance to SQL Server on a Microsoft Azure virtual machine. The instance has 30 databased that consume a total of 2 TB of disk space.



The instance sustains more than 30,000 transactions per second.

You need to provision storage for the virtual machine. The storage must be able to support the same load as the on-premises deployment.

Solution: You create one storage account that has one container. You create multiple VHDs in the container. Does this meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

Each Storage Account handles up to 20,000 IOPS, and 500TB of data.

References: <https://www.tech-coffee.net/understand-microsoft-azure-storage-for-virtual-machines/>

**NEW QUESTION 54**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that contains a financial database hosted on a storage area network (SAN).

The financial database has the following characteristics:

The database is continually modified by users during business hours from Monday through Friday between 09:00 hours and 17:00 hours. Five percent of the existing data is modified each day.

The Finance department loads large CSV files into a number of tables each business day at 11:15 hours and 15:15 hours by using the BCP or BULK INSERT commands. Each data load adds 3 GB of data to the database.

These data load operations must occur in the minimum amount of time.

A full database backup is performed every Sunday at 10:00 hours. Backup operations will be performed every two hours (11:00, 13:00, 15:00, and 17:00) during business hours.

You need to ensure that the backup size is as small as possible. Which backup should you perform every two hours?

- A. NORECOVERY
- B. FULL
- C. NO\_CHECKSUM
- D. CHECKSUM
- E. Differential
- F. BULK\_LOGGED
- G. STANDBY
- H. RESTART
- I. SKIP
- J. Transaction log
- K. DBO ONLY
- L. COPY\_ONLY
- M. SIMPLE
- N. CONTINUE AFTER ERROR

**Answer: J**

**Explanation:**

Minimally, you must have created at least one full backup before you can create any log backups. After that, the transaction log can be backed up at any time unless the log is already being backed up.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/transaction-log-backups-sql-server>

**NEW QUESTION 55**

- (Exam Topic 7)

You administer two instances of Microsoft SQL Server 2014. You deploy an application that uses a database on the named instance.

The application is unable to connect to the database on the named instance. You need to ensure that the application can connect to the named instance. What should you do?

- A. Configure the application as data-tiered.
- B. Open port 1433 on the Windows firewall on the server.
- C. Configure the named SQL Server instance to use an account that is a member of the Domain Admins group.
- D. Start the SQL Server Browser Service.

**Answer: D**

**Explanation:**

The SQL Server Browser program runs as a Windows service. SQL Server Browser listens for incoming requests for Microsoft SQL Server resources and provides information about SQL Server instances installed on the computer. SQL Server Browser contributes to the following actions:

References: [https://technet.microsoft.com/en-us/library/ms181087\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms181087(v=sql.105).aspx)

**NEW QUESTION 60**

- (Exam Topic 7)

You administer a Windows Azure SQL Database database named Human\_Resources. The database contains 2 tables named Employees and SalaryDetails. You add two Windows groups as logins for the server:

You need to grant users access according to the following requirements: What should you do?

- A. Create a database role called Employees. Add CORP\Employees to the db\_datareader rol
- B. Add all company employees except HR administrators to the Employees rol
- C. Deny SELECT access to the SalaryDetails table to the Employees role.
- D. Create a database role called HRAdmins. Add all company employees except HR administrators to the db\_datareader rol
- E. Add all HR administrators to the HRAdmins rol
- F. Grant SELECT access to the SalaryDetails table to the HRAdmins role. Deny SELECT access to the SalaryDetails table to the db\_datareader role.
- G. Create two database roles: Employees and HRAdmin

- H. Add all company employees to the Employees role. Add HR administrators to the HRAdmins rol
- I. Grant SELECT access to all tables except SalaryDetails to the Employees rol
- J. Grant SELECT access to the SalaryDetails table to the HRAdmins rol
- K. Deny SELECT access to the SalaryDetails table to the Employees role.
- L. Create a database role called Employees. Add all HR administrators to the db\_datareader rol
- M. Add all company employees to the Employees rol
- N. Grant SELECT access to all tables except the SalaryDetails table to the Employees rol
- O. Deny SELECT access to the SalaryDetails table to the Employees role.

**Answer:** D

**Explanation:**

Members of the db\_datareader fixed database role can run a SELECT statement against any table or view in the database.

References: [https://technet.microsoft.com/en-us/library/ms188629\(v=sql.90\).aspx](https://technet.microsoft.com/en-us/library/ms188629(v=sql.90).aspx)

**NEW QUESTION 61**

- (Exam Topic 7)

You are the administrator of a Microsoft SQL Server 2014 server.

Some applications consume significant resources. You need to manage the server workload by restricting resource-intensive applications

You need to dynamically limit resource consumption. What should you do?

- A. Configure Resource Pools, Workload Groups, and Classifier Function, and then enable the Resource Governor
- B. Set up Service Broker to ensure that application are not allowed to consume more than the specified amount of resource
- C. Create a new rule for each application that sets the resource limit allowed
- D. Create a new plan Guide with a Scope Type of sql and define the resource limits for each application

**Answer:** A

**Explanation:**

In the SQL Server Resource Governor, a resource pool represents a subset of the physical resources of an instance of the Database Engine. Resource Governor enables you to specify limits on the amount of CPU, physical IO, and memory that incoming application requests can use within the resource pool. Each resource pool can contain one or more workload groups. When a session is started, the Resource Governor classifier assigns the session to a specific workload group, and the session must run using the resources assigned to the workload group.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor-resou>

**NEW QUESTION 63**

- (Exam Topic 7)

You have Microsoft SQL Server on a Microsoft Azure Virtual machine that has a 4-TB database.

You plan to configure daily backups for the database. A single full backup will be approximately 1.5 TB of compressed data.

You need to ensure that the last backups are retained. Where should you store the daily backups?

- A. Local storage
- B. Page blob storage
- C. Virtual disks
- D. Block blob storage.

**Answer:** D

**Explanation:**

When backing up to Microsoft Azure blob storage, SQL Server 2016 supports backing up to multiple blobs to enable backing up large databases, up to a maximum of 12.8 TB. This is done through Block Blobs.

References:

**NEW QUESTION 68**

- (Exam Topic 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 master database.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

The master database does not handle alerts and jobs. It records all the system-level information for a SQL Server system. This includes instance-wide metadata such as logon accounts, endpoints, linked servers, and system configuration settings.

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>

**NEW QUESTION 70**

- (Exam Topic 7)

You are designing a Windows Azure SQL Database for an order fulfillment system. You create a table named Sales.Orders with the following script.

```
CREATE TABLE Sales.Orders
(
    OrderID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
    OrderDate datetimeoffset NOT NULL,
    CustomerID int NOT NULL
);
```

Each order is tracked by using one of the following statuses:

Fulfilled  
Shipped  
Ordered  
Received

You need to design the database to ensure that that you can retrieve the following information:

The current status of an order  
The previous status of an order.  
The date when the status changed.  
The solution must minimize storage.

More than one answer choice may achieve the goal. Select the BEST answer.

- A. To the Sales.Orders table, add three columns named Status, PreviousStatus and ChangeDat
- B. Update rows as the order status changes.
- C. Create a new table named Sales.OrderStatus that contains three columns named OrderID, StatusDate, and Statu
- D. Insert new rows into the table as the order status changes.
- E. Implement change data capture on the Sales.Orders table.
- F. To the Sales.Orders table, add three columns named FulfilledDate, ShippedDate, and ReceivedDate.Update the value of each column from null to the appropriate date as the order status changes.

**Answer:** A

**Explanation:**

This stores only the minimal information required.

#### NEW QUESTION 71

- (Exam Topic 7)

You are a database developer for an application hosted on a Microsoft SQL Server 2014 server. The database contains two tables that have the following definitions:

```
CREATE TABLE Customer
(CustomerID int NOT NULL PRIMARY KEY,
 CustomerName varchar(50) NOT NULL)

CREATE TABLE Orders
(OrderID int NOT NULL PRIMARY KEY,
 CustomerID int NOT NULL FOREIGN KEY REFERENCES Customer (CustomerID),
 OrderAmount money NOT NULL,
 ShippingCountry varchar(50) NOT NULL)
```

Global customers place orders from several countries. You need to view the country from which each customer has placed the most orders. Which Transact-SQL query do you use?

- A. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,RANK() OVER (PARTITION BY CustomerIDORDER BY COUNT(OrderAmount) DESC) AS RnkFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDWHERE o.Rnk = 1
- B. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM(SELECT c.CustomerID, c.CustomerName, o.ShippingCountry, RANK()OVER (PARTITION BY CustomerIDORDER BY COUNT(o.OrderAmount) ASC) AS RnkFROM Customer cINNER JOIN Orders oON c.CustomerID = o.CustomerIDGROUP BY c.CustomerID, c.CustomerName, o.ShippingCountry) cs WHERE Rnk = 1
- C. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,RANK() OVER (PARTITION BY CustomerIDORDER BY OrderAmount DESC) AS RnkFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDWHERE o.Rnk = 1
- D. SELECT c.CustomerID, c.CustomerName, o.ShippingCountry FROM Customer cINNER JOIN(SELECT CustomerID, ShippingCountry,COUNT(OrderAmount) DESC) AS OrderAmountFROM OrdersGROUP BY CustomerID, ShippingCountry) AS oON c.CustomerID = o.CustomerIDORDER BY OrderAmount DESC

**Answer:** A

**Explanation:**

Use descending (DESC) ordering.

To order by the number of orders we use ORDER BY COUNT(OrderAmount). Finally a WHERE close is needed: WHERE o.Rnk = 1

#### NEW QUESTION 76

- (Exam Topic 7)

You manage an on-premises, multi-tier application that has the following configuration:

Two SQL Server 2012 databases named SQL1 and SQL2

Two application servers named AppServer1 and AppServer2 that run IIS You plan to move your application to Azure.

You need to ensure that during an Azure update cycle or a hardware failure, the application remains available.

Which two deployment configurations should you implement? Each correct answer presents part of the solution.



- A. Deploy AppServer1 and AppServer2 in a single availability set.
- B. Deploy all servers in a single availability set.
- C. Deploy SQL1 and AppServer1 in a single availability set.
- D. Deploy SQL2 and AppServer2 in a single availability set.
- E. Deploy SQL1 and SQL2 in a single availability set.

**Answer:** AE

**Explanation:**

You should deploy AppServer1 and AppServer2 in a single availability set. You should deploy SQL1 and SQL2 in a single availability set.

Note: Using availability sets allows you to build in redundancy for your Azure services. By grouping related virtual machines and services (tiers) into an availability set (in this case, deploying both of your databases into an availability set), you ensure that if there is a planned or unplanned outage, your services will remain available. At the most basic level, virtual machines in an availability set are put into a different fault domain and update domain. An update domain allows virtual machines to have updates installed and then the virtual machines are rebooted together.

If you have two virtual machines in an availability set, each in its own update domain, a rebooting of one server does not bring down all of the servers in a given tier. A fault domain operates in the same manner, so if there is a physical problem with a server, rack, network, or other service, both machines are separated, and services will continue.

**NEW QUESTION 79**

- (Exam Topic 7)

A company has an on-premises Microsoft SQL Server 2017 infrastructure. The storage area network (SAN) that supports the SQL infrastructure has reached maximum capacity.

You need to recommend a solution to reduce on-premises storage use without changing the application. What should you do?

- A. Configure an Express Route connection to Microsoft Azure.
- B. Configure a Microsoft Azure Key Vault.
- C. Configure geo-replication on the SAN.
- D. Configure SQL Server Stretch Database in Microsoft Azure.

**Answer:** D

**Explanation:**

Stretch warm and cold transactional data dynamically from SQL Server to Microsoft Azure with SQL Server Stretch Database. Unlike typical cold data storage, your data is always online and available to query. Benefit from the low cost of Azure rather than scaling expensive, on-premises storage.

References:

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

**NEW QUESTION 84**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2016 instance.

You need to configure a new database to support FILETABLES. What should you do? Choose all that apply.

- A. Disable FILESTREAM on the Database.
- B. Enable FILESTREAM on the Server Instance.
- C. Configure the Database for Partial Containment.
- D. Create a non-empty FILESTREAM file group.
- E. Enable Contained Databases on the Server Instance.
- F. Set the FILESTREAM directory name on the Database.

**Answer:** BDF

**Explanation:**

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/blob/enable-the-prerequisites-for-filetable>

**NEW QUESTION 85**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 database instance.

You plan to migrate the database to Windows Azure SQL Database. You verify that all objects contained in the database are compatible with Windows Azure SQL Database.

You need to ensure that database users and required server logins are migrated to Windows Azure SQL Database.

What should you do?

- A. Use the copy database wizard
- B. Use the Database Transfer wizard
- C. Use SQL Server Management Studio to deploy the database to Windows Azure SQL Database
- D. Backup the database from the local server and restore it to Windows Azure SQL Database

**Answer:** C

**Explanation:**

You would need to use either the SQL Server Management Studio or Transact-SQL.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-cloud-migrate>

**NEW QUESTION 89**

- (Exam Topic 7)

You administer two Microsoft SQL Server 2014 servers. Each server resides in a different, untrusted domain. You plan to configure database mirroring.

You need to be able to create database mirroring endpoints on both servers. What should you do?

- A. Configure the SQL Server service account to use Network Service.
- B. Use a server certificate.
- C. Use a database certificate.
- D. Configure the SQL Server service account to use Local System.

**Answer:** B

**Explanation:**

To enable certificate authentication for database mirroring on a given server instance, the system administrator must configure each server instance to use certificates on both outbound and inbound connections.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/database-mirroring/use-certificates-for-a-database-mirrorin>

**NEW QUESTION 90**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance.

The instance contains a database that supports a retail sales application. The application generates hundreds of transactions per second and is online 24 hours per day and 7 days per week.

You plan to define a backup strategy for the database. You need to ensure that the following requirements are met:

No more than 5 minutes worth of transactions are lost. Data can be recovered by using the minimum amount of administrative effort.

What should you do? Choose all that apply.

- A. Configure the database to use the SIMPLE recovery model.
- B. Create a DIFFERENTIAL database backup every 4 hours.
- C. Create a LOG backup every 5 minutes.
- D. Configure the database to use the FULL recovery model.
- E. Create a FULL database backup every 24 hours.
- F. Create a DIFFERENTIAL database backup every 24 hours.

**Answer:** BCDE

**Explanation:**

The full recovery model uses log backups to prevent data loss in the broadest range of failure scenarios, and backing and restoring the transaction log (log backups) is required. The advantage of using log backups is that they let you restore a database to any point of time that is contained within a log backup (point-in-time

recovery). You can use a series of log backups to roll a database forward to any point in time that is contained in one of the log backups. Be aware that to minimize your restore time, you can supplement each full backup with a series of differential backups of the same data.

References: [https://technet.microsoft.com/en-us/library/ms190217\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms190217(v=sql.105).aspx)

**NEW QUESTION 94**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance. After a routine shutdown, the drive that contains tempdb fails.

You need to be able to start the SQL Server. What should you do?

- A. Modify tempdb location in startup parameters.
- B. Start SQL Server in minimal configuration mode.
- C. Start SQL Server in single-user mode.
- D. Configure SQL Server to bypass Windows application logging.

**Answer:** B

**Explanation:**

If you have configuration problems that prevent the server from starting, you can start an instance of Microsoft SQL Server by using the minimal configuration startup option.

When you start an instance of SQL Server in minimal configuration mode, note the following: Only a single user can connect, and the CHECKPOINT process is not executed.

Remote access and read-ahead are disabled. Startup stored procedures do not run.

tempdb is configured at the smallest possible size.

References:

<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/start-sql-server-with-minimal-configur>

**NEW QUESTION 97**

- (Exam Topic 7)

You are the database administrator for your company. Your company has one main office and two branch offices. You plan to create three databases named DB1, DB2, and DB3 that will be hosted on one Azure SQL Database server. You have the following requirements:

The main office must be able to connect to all three databases.

The branch offices must be able to connect to DB2 and DB3.

The branch offices must not be able to access DB1.

You need to configure transparent data encryption (TDE) for DB1. Which two actions should you perform? Each correct answer presents part of the solution.

- A. Run CREATE CERTIFICATE cert1 WITH Subject = TDE Cert1 on DB1.
- B. Connect to DB1.
- C. Run ALTER DATABASE DB1 SET ENCRYPTION ON;
- D. Connect to the master database.
- E. Run CREATE MASTER KEY on the master database.

**Answer:** BC

**Explanation:**

You should connect to DB1. To encrypt DB1, you connect directly to DB1. When you connect to DB1. You use your dbmanager or administrative credentials.

You should run ALTER DATABASE DB1 SET ENCRYPTION ON.

You use the ALTER DATABASE DB1 SET ENCRYPTION ON statement to encrypt the database. This is the statement that turns on TDE for Azure SQL Database.

**NEW QUESTION 101**

- (Exam Topic 7)

You administer a Microsoft SQL Server 2014 instance that has multiple databases. You have a two-node SQL Server failover cluster. The cluster uses a storage area network (SAN). You discover I/O issues. The SAN is at capacity and additional disks cannot be added.

You need to reduce the I/O workload on the SAN at a minimal cost. What should you do?

- A. Move user databases to a local disk.
- B. Expand the tempdb data and log files
- C. Modify application code to use table variables
- D. Move the tempdb files to a local disk

**Answer:** D

**Explanation:**

The use of local disks for TempDB allows us to have more flexibility when configuring for optimal performance. It is a common performance recommendation to create the TempDB database on the fastest storage available. With the capability to utilize local disk for TempDB placement we can easily utilize disks that are larger, have a higher rotational speed or use SSD disks.

References: <https://www.mssqltips.com/sqlservertip/2817/sql-server-2012-cluster-with-tempdb-on-local-disk/>

**NEW QUESTION 105**

- (Exam Topic 7)

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 have Microsoft SQL Server on a Microsoft Azure virtual machine that has a database named DB1. You discover that DB1 experiences WRITE\_LOG waits that are longer than 50 ms.

You need to reduce the WRITE\_LOG wait time. Solution: Move the transaction logs to a faster disk. Does this meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:**

Section: Deploy and migrate applications

In SQL Server, if we have a transactional based system and find a high WRITELOG wait type this is a performance bottleneck and can cause the transaction log file to grow rapidly and frequently.

It is being recommended to SQL server users that they must archive the log files on a separate disk for getting better performance.

References: <https://atdhebuja.wordpress.com/2016/06/20/resolving-sql-server-transaction-log-waits/>

**NEW QUESTION 110**

.....



## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your 70-765 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/70-765-dumps.html>