

Microsoft.Premium.70-473.by.VCEplus.170q

Number: 70-473 VCEplus
Passing Score: 800
Time Limit: 120 min
File Version: 5.4



Exam Code: 70-473

Exam Name: Designing and Implementing Cloud Data Platform Solutions

Certification Provider: Microsoft

Corresponding Certification: Microsoft Specialist

Website: www.vceplus.com

Free Exam: <https://vceplus.com/exam-70-473/>

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in 70-473 exam products and you get latest questions. We strive to deliver the best 70-473 exam product for top grades in your first attempt.

Website: <https://vceplus.com>

VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Proseware, Inc.

QUESTION 1

Overview

General Overview

Proseware, Inc. is a software engineering company that has 100 employees. Proseware has sales, marketing, accounts, human resources IT, and development departments.

The IT department has one team dedicated to managing the internal resources and one team dedicated to managing customer resources, which are located in the company's hosting environment.

Proseware develops websites, basic web apps, and custom web apps. The websites and the apps are hosted and maintained in the hosting environment of Proseware.

Physical Locations

Proseware has two offices located in Seattle and Montreal. The Seattle office contains all of the hardware required to host its customers' websites, web apps, and databases. The Seattle office contains the IT team for the hosting environment.

The Montreal office contains all of the hardware required to host the company's internal applications, databases, and websites.

Each office connects directly to the Internet. Testing reveals that the minimum latency from the offices to Microsoft Azure is 20 ms.

Existing Environment

Internal Microsoft SQL Server Environment

Proseware uses a custom customer relationship management (CRM) application.

The internal Microsoft SQL Server environment contains two physical servers named CRM-A and CRM-B. Both servers run SQL Server 2012 Standard and host databases for the CRM application.

CRM-A hosts the principal instance and CRM-B hosts the mirrored instance of the CRM database. CRM-A also hosts databases for several other applications that are used by the company's internal applications.

CRM-A has a quad core processor and 12 GB of RAM. CRM-B has a dual core processor and 8 GB of RAM.

Custom Web Applications Environment

Some Proseware customers request custom web-based applications that require more than just databases, such as SQL Server Integration Services (SSIS) and CLR stored procedures.

Proseware uses a Hyper-V server named Host1. Host1 has four instances of SQL Server 2014 Enterprise in the host operating system. The instances are mirrored on a server named Host2.

Host1 also hosts four virtual machines named VM1, VM2, VM3, and VM4. VM1 has SQL Server 2005 Standard installed. VM2 has SQL Server 2005 Enterprise Edition installed. VM3 has SQL Server 2008 Standard Edition installed. VM4 has SQL Server 2008 R2 Standard Edition installed.

Host1 uses a SAN to store all of the data and log files for the four SQL Server instances and the four virtual machines.

Websites and Basic Web Apps Environment

Proseware has two physical servers named WebServer1 and WebData1. WebServer1 hosts basic web apps and websites for its customers. WebData1 has a database for each website and each basic web app that Proseware hosts. WebData1 has four cores and 8 GB of RAM.

Each website database contains customer information for billing purposes. Proseware generates a consolidated report that contains data from all of these databases.

The relevant databases on WebData1 are:

- CWDB: Currently 60 GB and is not expected to exceed 100 GB. CWDB contains a table named PersonalInfo.
- MovieReviewDB: Currently 5 GB and is not expected to exceed 10 GB.

Marketing Department

Proseware has a web app for the marketing department. The web app uses an Azure SQL database. Managers in the marketing department occasionally bulk load data by using a custom application. The database is updated daily.

Problem Statements

Proseware identifies the following issues:

- Lack of planning and knowledge has complicated the database environment.
- Customers who have web apps hosted on WebServer1 report frequent outages caused by failures on WebData1. The current uptime is less than 90 percent.
- Internally, users complain of slow performance by the CRM application when the databases fail over to CRM-B.
- WebData1 has no high-availability option for the databases or the server.
- An internal licensing audit of SQL Server identifies that Proseware is non-compliant. Host1, CRM-A, and CRM-B are licensed properly. VM1, VM2, VM3, VM4, and WebData1 are unlicensed.

Business Requirements

Proseware identifies the following business requirements:

- Upgrade the infrastructure to address the issues reported by the internal users and customers.
- Minimize upgrade costs associated with purchasing hardware and software.
- Ensure that all software is licensed properly.
- Minimize the complexity of the database environment.
- Consolidate the instances of SQL Server that support the custom web app environment.
- Implement a service level agreement (SLA) of 99.95 percent uptime for the website and basic web app environment.
- Implement a disaster recovery environment in Azure for the CRM application.
- Ensure that any changes to the SQL Server environments either maintain or increase overall performance.
- Migrate all web front ends to Azure.
- Reuse licenses, whenever possible.

- Minimize the administrative effort required to generate the internal reports from the website databases.

Security Requirements

Proseware hosts a database for a company named Contoso, Ltd. Currently, all of the employees at Contoso can access all of the data in the database.

Contoso plans to limit user access to the CWDU database so that customer service representatives can see only the data from the PersonalInfo table that relates to their own customers.

A.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

You are planning the consolidation of the databases from Host1.

You need to identify which methods to use to back up the data after the consolidation completes. What are two possible methods that achieve this goal? Each correct answer presents a complete solution.

- A. BACKUP TO URL
- B. AlwaysOn failover clustering
- C. a maintenanceplan
- D. AlwaysOn Availability Groups

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

QUESTION 3

You need to recommend a disaster recovery solution for the CRM application that meets the business requirements.

What should you recommend?

- A. backup and restore by using Windows Azure Storage

- B. log shipping
- C. AlwaysOn Availability Groups
- D. database mirroring

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 4

You are evaluating moving the data from WebData1 to an Azure SQL database.

You need to recommend a solution to generate the consolidated report for billing. The solution must meet the business requirements.

What should you include in the recommendation?

- A. SQL Server Integration Services (SSIS)
- B. SQL Server Analysis Services (SSAS)
- C. an elastic database job
- D. an elastic query



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 5

You are evaluating whether to replace CRM-B with an Azure virtual machine.

You need to identify the required virtual machine service tier to replace CRM-B. The solution must meet the following requirements:

- Ensure that the database is stored in premium storage.
- Meet the business requirements.

Which virtual machine service tier should you identify?

- A. Standard DS3
- B. Standard A6
- C. Standard GS2

D. Standard D3

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References: <https://azure.microsoft.com/en-gb/documentation/articles/virtual-machines-windows-sizes/>

QUESTION 6

HOTSPOT

You need to identify which methods to use to migrate MovieReviewDB and CWDB.

Which method should you identify for each database? To answer, select the appropriate options in the answer area.

Hot Area:

Hot Area:



Answer Area

CWDB:

	▼
Azure Import and Export Service	
Azure ExpressRoute	
the bcp utility	

MovieReviewDB:

	▼
Azure ExpressRoute	
the BACKUP and RESTORE statements	
the Deploy Database to Microsoft Azure Database task	

Correct Answer:

Answer Area

CWDB:

	▼
Azure Import and Export Service	
Azure ExpressRoute	
the bcp utility	

MovieReviewDB:

	▼
Azure ExpressRoute	
the BACKUP and RESTORE statements	
the Deploy Database to Microsoft Azure Database task	

Section: (none)

Explanation

Explanation/Reference:

References: <https://azure.microsoft.com/en-gb/documentation/articles/sql-database-cloud-migrate-compatible-using-ssms-migration-wizard/>

QUESTION 7

HOTSPOT

You are evaluating the migration of the databases from Host1 and WebData1 to Azure.

You need to recommend the most cost-effective solution for storing the database in Azure. The solution must meet the business requirements.

In the table below, recommend the most cost-effective storage solution for Host1 and WebData1. NOTE: Make only one selection in each column.

Hot Area:

Hot Area:

Answer Area

Solutions	Host1	WebData1
SQL Server 2014 Standard edition installed on an Azure virtual machine	<input type="radio"/>	<input type="radio"/>
SQL Server 2014 Enterprise edition installed on an Azure virtual machine	<input type="radio"/>	<input type="radio"/>
a single Azure SQL database on the Basic service tier	<input type="radio"/>	<input type="radio"/>
an Azure SQL Database elastic database pool on the Basic service tier	<input type="radio"/>	<input type="radio"/>
a single Azure SQL database on the Standard service tier	<input type="radio"/>	<input type="radio"/>
an Azure SQL Database elastic database pool on the Standard service tier	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Solutions	Host1	WebData1
SQL Server 2014 Standard edition installed on an Azure virtual machine	<input type="radio"/>	<input type="radio"/>
SQL Server 2014 Enterprise edition installed on an Azure virtual machine	<input checked="" type="radio"/>	<input type="radio"/>
a single Azure SQL database on the Basic service tier	<input type="radio"/>	<input type="radio"/>
an Azure SQL Database elastic database pool on the Basic service tier	<input type="radio"/>	<input type="radio"/>
a single Azure SQL database on the Standard service tier	<input type="radio"/>	<input type="radio"/>
an Azure SQL Database elastic database pool on the Standard service tier	<input type="radio"/>	<input checked="" type="radio"/>

Section: (none)

Explanation

Explanation/Reference:

QUESTION 8

DRAG DROP

You plan to implement row-level security for the CWDB database.

You create the fn_limitusers function under the restriction schema.

You need to create the policy.

How should you complete the policy? To answer, drag the appropriate elements to the correct locations. Each element 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:

Select and Place:

Code elements

ADD BLOCK PREDICATE restriction.fn_limitusers (CRS)

ADD FILTER PREDICATE restriction.fn_limitusers (CRS)

ON CWDB

ON PersonalInfo

WITH (STATE = ON)

WITH SCHEMABINDING

Answer area

CREATE SECURITY POLICY CSRFILTER

Code element

Code element

Code element

Correct Answer:

Code elements

ADD BLOCK PREDICATE restriction.fn_limitusers (CRS)

ON CWDB

WITH SCHEMABINDING

Answer area

CREATE SECURITY POLICY CSRFILTER

ADD FILTER PREDICATE restriction.fn_limitusers (CRS)

ON PersonalInfo

WITH (STATE = ON)

Section: (none)

Explanation

Explanation/Reference:

QUESTION 9

DRAG DROP

A marketing department manager reports that the marketing department database contains incorrect data. The manager reports that the data was correct yesterday.

You need to recommend a method to recover the data.

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

Select and Place:

Select and Place:

Actions

Answer Area

Perform a point-in-time restore.

Modify the firewall rules.

Rename the database.

Enable Change Tracking.

Create a new Azure SQL database server.

Delete the database.

Migrate the database to an on-premises instance.



Correct Answer:

Actions

Modify the firewall rules.

Enable Change Tracking.

Create a new Azure SQL database server.

Migrate the database to an on-premises instance.

Answer Area

Rename the database.

Perform a point-in-time restore.

Delete the database.



Section: (none)

Explanation

Explanation/Reference:

Explanation:

Rename the current live database.

Restore a backup of the database as a new database with the same name as the original database. Now you can delete the old database.

ADatum Corporation

QUESTION 1

Overview

General Overview

ADatum Corporation is a real estate firm that has offices throughout North America.

ADatum has a main office and four branch offices. The main office is located in Seattle. The branch offices are located in New York, Montreal, Denver, and Vancouver.

Existing Environment

Network Infrastructure

The network contains one Active Directory domain named adatum.com. Each office contains one domain controller.

Each office has a 100-Mbps connection to the Internet that is 20 percent saturated on average. The offices connect to each other through the Internet by using VPN appliances.

ADatum uses the public IP addresses shown in the following table.

Office	Public IP address
Seattle	131.107.1.6
New York	131.107.2.6
Montreal	131.107.3.6
Denver	131.107.4.6
Vancouver	131.107.5.6

SQL Server Infrastructure

In the main office, ADatum hosts a Microsoft SQL Server instance on a server named SQL1. SQL1 has a 400-GB database named Listings. Log-shipped copies of Listings are present in each branch office. The copies are used for reporting. Currently, all of the SQL Server instances run SQL Server 2014 Enterprise edition.

Each branch office has an application server that hosts an application named App1. App1 is configured to connect to the Listings database on each local SQL Server instance for reporting and to connect to the SQL Server instance in the main office for any updates to property listings.

The main office also has an application server that hosts App1. The application server connects to the local Listings database for reporting and for any updates to the property listings.

Historic activity of the Listings database shows a maximum of 475 concurrent requests from as many as 200 concurrent connections.

User Issues

Users report that, frequently, they are disconnected from the Listings database when they run reports. Users also report that there is an unacceptable delay between when a property listing is updated and when the updated listing appears in the listings reports.

Developers report concerns about the lack of a testing environment in which code changes can be validated before being deployed to the production Listings database.

Requirements

Business Requirements

ADatum identifies the following business requirements:

- Minimize costs, whenever possible.
- Ensure that confidential data is encrypted at all times.
- Ensure that the primary database is hosted in Microsoft Azure.
- Ensure that all production databases maintain 99.9 percent availability.
- Ensure that all of the data between the offices and Azure is encrypted.

Planned changes

ADatum plans to implement the following changes:

- Move the primary database to Azure.
- Implement a data warehouse for reporting to offload reporting from the transactional Listings database.

Technical Requirements

ADatum identifies the following technical requirements:

- A test environment that has a 200-GB subset of data from the Listings database must be implemented. The new database will be named ListTest. The new test environment will have a maximum of 10 concurrent connections.
- The migration of the Listings database must be completed in less than 60 minutes. During the migration, data must be prevented from being modified.
- The firewall settings of the Azure SQL databases must be configured to provide access to the main office only.
- Changes to the settings and the properties of the Listings database must be audited at all times.
- Access to the Clients table must be audited and data from the audit must be queryable.
- The query performance of the ListTest database must be monitored at all times.
- Reporting must be offloaded from the transactional Listings database.

High-Availability and Recovery Requirements

ADatum identifies the following high-availability and recovery requirements:

- All production databases must support automatic failover.

- Backups for the Listings database must be stored in Azure.
- The database must be recoverable if a major data loss occurs.
- A weekly backup of the ListTest database must be maintained on-premises.
- SQL1 must be integrated into the high-availability solution as a reporting server.

A.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

You need to recommend a backup solution for the ListTest database. What should you include in the recommendation?

- A. Extract a data-tier application (DAC).
- B. Use the bcp command.
- C. Use the SQL Server Migration Assistant (SSMA).
- D. Export a data-tier application (DAC).



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 3

You need to recommend a solution to migrate the Listings database to the cloud.
What should you recommend?

- A. Stage the Listings database on an Azure virtual machine prior to the outage. During the outage, perform a data import from the main office by using SQLServer Management Studio.
- B. Implement log-shipping between the main office and an Azure virtual machine prior to the migration date. During the planned outage, perform a final logbackup, restore the backup to the secondary, and then switch the secondary to the primary role.
- C. Run a full backup during the outage and restore the backup to the Azure virtual machine.
- D. Implement merge replication between the main office and the Azure virtual machine.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 4

You are evaluating whether an Azure SQL Database elastic database pool suits your workload and usage patterns.

What are two possible ways to identify the elastic database transaction units (eDTUs)? Each correct answer presents a complete solution.

- A. Aggregate data from sys.dm_os_wait_stats.
- B. Run the Database Engine Tuning Advisor.
- C. Run the Service Tier Advisor.
- D. Aggregate data from sys.dm_db_resource_stats.
- E. Aggregate data from sys.dm_os_performance_counters.

Correct Answer: DE

Section: (none)

Explanation



Explanation/Reference:

QUESTION 5

You need to recommend a solution to migrate the Listings database to the cloud.

What should you recommend?

- A. Stage the Listings database on an Azure virtual machine prior to the outage. During the outage, perform a data import from the main office by using SQLServer Management Studio.
- B. Enable mirroring between the main office and the Azure virtual machine. During the planned outage, change the mirroring roles on the Azure virtual machine to become the primary server, and then break the mirror.
- C. Run a full backup during the outage and restore the backup to the Azure virtual machine.
- D. Implement merge replication between the main office and the Azure virtual machine.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

<https://blogs.msdn.microsoft.com/buckwoody/2013/01/08/microsoft-windows-azure-disaster-recovery-options-for-on-premises-sql-server/>

QUESTION 6

DRAG DROP

You need to configure auditing to meet the technical requirements.

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.

Select and Place:

Actions

Create and start a server audit.

Configure a database audit specification.

Enable auditing for failed and successful logins.

Read the audit file.

Enable C2 audit mode.

Answer Area



Correct Answer:

Actions

Read the audit file.

Enable C2 audit mode.



Answer Area

Create and start a server audit.

Enable auditing for failed and successful logins.

Configure a database audit specification.



Section: (none)

Explanation

Explanation/Reference:

References: <http://solutioncenter.apexsql.com/how-to-setup-and-use-sql-server-audit-feature/>

QUESTION 7

DRAG DROP

You need to implement the monitoring solution for the Listings database.

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.

Select and Place:

Actions

Answer Area

Use the **sys.fn_get_audit_file()** function to review and report on the trapped events.

Execute the following statement.

```
CREATE EVENT SESSION
ListingsDBChange ON DATABASE
ADD EVENT
sqlserver.object_altered ( SET
collect_database_name = (1)
ACTION ( sqlserver.sql_text
,sqlserver.nt_username
,sqlserver.server_principal_name
,sqlserver.client_hostname )
WHERE object_type =
'DATABASE'
AND
sqlserver.database_name =
'Listings'
)
ADD TARGET package0.event_file
( SET filename =
N'C:\Database\XE\ListingsDBChange
WITH (STARTUP_STATE = ON);
```

Execute the following statement.

```
ALTER EVENT SESSION
ListingsDBChange
ON SERVER
STATE = START;
GO
```

Use the
sys.fn_xe_file_target_read_file()
function to review and report on the trapped
events.

Execute the following statement

Correct Answer:



Actions

Use the **sys.fn_get_audit_file()** function to review and report on the trapped events.

Execute the following statement.

```
CREATE EVENT SESSION
ListingsDBChange ON DATABASE
ADD EVENT
sqlserver.object_altered ( SET
collect_database_name = (1)
ACTION ( sqlserver.sql_text
,sqlserver.nt_username
,sqlserver.server_principal_name
,sqlserver.client_hostname )
WHERE object_type =
'DATABASE'
AND
sqlserver.database_name =
'Listings'
)
ADD TARGET package0.event_file
( SET filename =
N'C:\Database\XE\ListingsDBChange
WITH (STARTUP_STATE = ON);
```

Answer Area

Execute the following statement.

```
CREATE EVENT SESSION
ListingsDBChange ON SERVER
ADD EVENT
sqlserver.object_altered ( SET
collect_database_name = (1)
ACTION ( sqlserver.sql_text
,sqlserver.nt_username
,sqlserver.server_principal_name
,sqlserver.client_hostname )
WHERE object_type =
'DATABASE'
AND
sqlserver.database_name =
'Listings'
)
ADD TARGET package0.event_file
( SET filename =
N'C:\Database\XE\ListingsDBChange
WITH (STARTUP_STATE = ON);
```

Execute the following statement.

```
ALTER EVENT SESSION
ListingsDBChange
ON SERVER
STATE = START;
GO
```

Use the
sys.fn_xe_file_target_read_file()
function to review and report on the trapped
events.

Section: (none)

Explanation

Explanation/Reference:





QUESTION 8

DRAG DROP

You need to back up the Listings database to meet the high-availability and recovery requirements.

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.

Select and Place:

Actions		Answer Area
Create a database snapshot.	 	 
Create storage keys.		
Create a credential.		
Back up the database to Azure Blob storage.		
Create a database copy.		

Correct Answer:

Actions

Create a database snapshot.

Create a database copy.



Answer Area

Create a credential.

Create storage keys.

Back up the database to Azure Blob storage.



Section: (none)

Explanation

Explanation/Reference:

QUESTION 9

DRAG DROP

You need to recommend a solution to implement high availability for the Listings database.

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

Select and Place:

Actions

Answer Area

Configure availability groups that have the primary replica on an Azure virtual machine and a secondary read-only replica in the main office.

Configure mirroring between the Azure virtual machine and the main office in high-safety mode that supports automatic failover.

Create a failover cluster between SQL1 and an Azure virtual machine.

Configure a SQL Server instance on an Azure virtual machine as the witness for database mirroring.

Configure availability groups between two Azure virtual machines in geographically dispersed zones. Place a secondary read-only replica on the second virtual machine.

Create a VPN tunnel between the main office and the Azure virtual machine.



Correct Answer:

Actions

Configure availability groups that have the primary replica on an Azure virtual machine and a secondary read-only replica in the main office.

Create a failover cluster between SQL1 and an Azure virtual machine.

Configure availability groups between two Azure virtual machines in geographically dispersed zones. Place a secondary read-only replica on the second virtual machine.

Answer Area

Create a VPN tunnel between the main office and the Azure virtual machine.

Configure mirroring between the Azure virtual machine and the main office in high-safety mode that supports automatic failover.

Configure a SQL Server instance on an Azure virtual machine as the witness for database mirroring.



Section: (none)

Explanation

Explanation/Reference:

QUESTION 10

Which three components are required to meet the encryption requirements for the ListTest database? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. a Column Master Key
- B. a certificate

- C. an Always Encrypted enabled driver
- D. a Column Encryption Key
- E. a service master key
- F. a database master key

Correct Answer: ACD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 11

DRAG DROP

You need to implement monitoring for the ListTest database.

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

Select and Place:

Select and Place:



Statements

```
ALTER EVENT SESSION XE_azuresqlldbListing
ON DATABASE
STATE=START;
GO
```

```
CREATE
EVENT SESSION XE_azuresqlldbListing
ON DATABASE
ADD EVENT
sqlserver.sql_statement_starting
(
ACTION(sqlserver.sql_text)
WHERE statement LIKE '%UPDATE Client%'
)
ADD TARGET
package0.ring_buffer
(SET
max_memory = 500 --Units of KB
);
GO
```

```
SELECT *
FROM sys.fn_xe_file_target_read_file
('XE_azuresqlldbListing',NULL, NULL, NULL)
```

```
SELECT CAST ([target_data] AS XML) AS
target_data
FROM sys.dm_xe_sessions AS xt
INNER JOIN sys.dm_xe_sessions AS xs
ON xs.address = xt.event_session_address
WHERE xs.name = N'XE_azuresqlldbListing'
AND xt.target_name = N'ring_buffer'
;
```

Answer Area

Correct Answer:



Statements

```
SELECT *
FROM sys.fn_xe_file_target_read_file
('XE_azuresqlldbListing',NULL, NULL, NULL)
```

Answer Area

```
CREATE
EVENT SESSION XE_azuresqlldbListing
ON DATABASE
ADD EVENT
sqlserver.sql_statement_starting
(
ACTION(sqlserver.sql_text)
WHERE statement LIKE '%UPDATE Client%'
)
ADD TARGET
package0.ring_buffer
(SET
max_memory = 500 --Units of KB
);
GO
```

```
ALTER EVENT SESSION XE_azuresqlldbListing
ON DATABASE
STATE=START;
GO
```

```
SELECT CAST ([target_data] AS XML) AS
target_data
FROM sys.dm_xe_sessions AS xt
INNER JOIN sys.dm_xe_sessions AS xs
ON xs.address = xt.event_session_address
WHERE xs.name = N'XE_azuresqlldbListing'
AND xt.target_name = N'ring_buffer'
;
```

Section: (none)

Explanation

Explanation/Reference:

QUESTION 12

DRAG DROP

You need to configure server-level access to the ListTest database. adatum

What statement should you use? To answer, drag the appropriate elements to the correct locations. Each element may be used 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.

Select and Place:



Code Elements

EXECUTE sp_set_firewall_rule

EXEC sp_set_database_firewall_rule

'131.107.1.6'

'131.107.5.6'

'127.0.0.1'

'255.255.255.0'

'0.0.0.0'

Answer Area

@name = N'ADatumDBAFirewallRule',

@start_ip_address =

@end_ip_address =



Correct Answer:

Code Elements

EXEC sp_set_database_firewall_rule

'127.0.0.1'

'255.255.255.0'

'0.0.0.0'

Answer Area

EXECUTE sp_set_firewall_rule

@name = N'ADatumDBAFirewallRule',

@start_ip_address =

'131.107.1.6'

@end_ip_address =

'131.107.5.6'



Section: (none)

Explanation

Explanation/Reference:

Contoso, Ltd.

QUESTION 1

Overview

Contoso, Ltd. is a national scientific research company that has sales, marketing, and research departments.

Contoso has a main office in Dallas and more than 20 satellite offices across the United States. Some employees work off-site at customer locations.

Contoso is expected to double in size during the next two years.

Existing Environment

The Microsoft SQL Server environment contains the servers configured as shown in the following table.

Server name	Role	Application name
SQL01	Online Transaction Processing (OLTP) primary database	One Sale
SQL02	A mirror of SQL01 and used for ad hoc reporting	One Report
SQL03	A data warehouse	One Data
Research01	A database server used by the research department	Research Gen II

Contoso currently uses SQL Server 2008 R2 on all of the database servers. SQL01 currently requires six cores and 50 GB of memory.

Once a day, database snapshots are taken on the mirror to provide ad hoc reporting and to load the data warehouse.

One report is used for ad hoc reporting by using the mirrored copy of the OLTP database.

Research Gen II is a Windows application that users currently access by using Remote Desktop.

Requirements

Business Goals

Contoso identifies the following business goals:

- Minimize the costs associated with purchasing hardware and software.
- Prevent any negative impact on performance when moving database platforms to the cloud.
- Audit user access to and query execution on all of the databases in the research department.

Planned changes

Contoso plans to implement the following changes:

- Upgrade SQL03 from SQL Server 2008 R2 to SQL Server 2014 on Windows Server 2012 R2.
- Migrate SQL01 and SQL02 to Infrastructure as a Service (IAAS) on SQL Server 2014.
- Leverage cloud services for a disaster recovery location and for high availability.
- Replace Research Gen II with a cloud-based application named One Research. The database for this application will be migrated to a cloud service.

Departmental Requirements

Senior research project leaders must be able to upload databases to the cloud. Initially, the databases will be less than 3 GB. Some databases may grow to 50 GB. The databases must be recoverable to any point during the past two weeks.

The research project leaders must monitor the performance of their databases. They must view statistics, such as query performance and overall database performance.

The marketing department plans to use cloud services to host web applications for marketing campaigns. The databases for the web applications have the requirements shown in the following table.

Marketing campaign	Database size	Number of users	Recovery point objective (RPO)	High-availability requirement
Campaign 1	1 GB	10	2 weeks	Failover
Campaign 2	200 GB	100	2 weeks	Failover and read-only
Campaign 3	200 GB	100	3 weeks	Failover

Technical Requirements

Contoso identifies the following technical requirements:

- OLTP will be offloaded to Microsoft Azure virtual machines.

- The data for One Report must not be more than 15 minutes old.
- All of the databases for the OLTP system must be implemented on solid state drives (SSDs).
- A solution for the OLTP system must be created to provide disaster recovery and reporting between SQL01 and SQL02.
- All databases, except for the research and marketing databases, must be recoverable to any point during the last 30 days.
- A solution for the data warehouse system must be created to provide disaster recovery between SQL03 and an Azure virtual machine named SQL04.

A.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

You need to recommend an auditing solution that meets the business goals.

Which three events should you include in the recommendation? Each correct answer presents part of the solution.

- A. Transaction Management
- B. Stored Procedure
- C. Plain SQL
- D. Login
- E. Failed Connections

Correct Answer: ABD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 3

You are designing the data warehouse.

You need to recommend a solution that meets the business goals and the technical requirements.

What should you recommend?

- A. Create an AlwaysOn availability group between SQL03 and SQL04.
- B. Host the data warehouse on an Azure virtual machine.
- C. Create an Azure HDInsight cluster.
- D. Create an Azure DocumentDB database.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 4

You are evaluating the use of active geo-replication for one of the research department databases.

You need to identify which service tier to use for the database. The solution must meet the department requirements and the business requirements.

What should you identify?

- A. Premium P1
- B. Premium P14
- C. Standard S0
- D. Basic



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-service-tiers/>

QUESTION 5

You need to ensure that the performance statistics for the research department are collected. What should you enable?

- A. Query Performance Insight

- B. Scale and Configuration
- C. Dynamic Data Masking
- D. Index advisor

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-query-performance/>

QUESTION 6

You need to tell the research project leaders how to migrate their databases.

Which task should you instruct the leaders to use from SQL Server Management Studio?

- A. Extract Data-tier Application
- B. Deploy Database to a Microsoft Azure VM
- C. Deploy Database to a Microsoft Azure SQL Database
- D. Copy Database



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

References:

<https://azure.microsoft.com/en-gb/documentation/articles/sql-database-cloud-migrate-compatible-using-ssms-migration-wizard/>

QUESTION 7

You establish a site-to-site VPN to Microsoft Azure.

You need to implement a high-availability solution for the data warehouse. The solution must meet the business goals and the technical requirements.

Which two actions should you perform? Each correct answer presents part of the solution.

- A. Configure the database to be in Full recovery model.
- B. Create an AlwaysOn availability group between SQL03 and SQL04.

- C. Implement iSCSI storage.
- D. Migrate SQL03 to an Azure virtual machine and configure an AlwaysOn availability group.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

QUESTION 8

DRAG DROP

You need to implement a storage architecture for the Azure virtual machines that will host the OLTP database. The solution must meet the business goals and technical requirements.

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.

Select and Place:

Actions

Create a single database file.

Assign a drive letter to each drive.

Use Storage Spaces to create a striped pool.

Create Azure virtual machines that have eight data drives.

Create tables in the file group.

Answer Area

>
 <

Correct Answer:

• • • • •

Actions

Assign a drive letter to each drive.

Create tables in the file group.



Answer Area

Create Azure virtual machines that have eight data drives.

Use Storage Spaces to create a striped pool.

Create a single database file.

Section: (none)

Explanation

Explanation/Reference:



QUESTION 9

You need to identify which service tier to use for the research department databases. The solution must meet the business goals and the department requirements.

What should you identify?

- A. an Azure virtual machine on the Basic service tier
- B. an Azure SQL database on the Basic service tier
- C. an Azure virtual machine on the Standard service tier
- D. an Azure SQL database on the Standard service tier
- E. an Azure SQL database on the Premium service tier

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 10

You are planning the migration of SQL01 to an Azure virtual machine.

You need to identify which service tier to use for the virtual machine. The solution must meet the business goals and the technical requirements.

Which virtual machine service tier should you identify?

- A. Standard D3
- B. Standard A7
- C. Standard G2
- D. Standard DS13

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



Relecloud

QUESTION 1

General overview

Relecloud is an international data services company that has 1000 users. The company has a sales department, a marketing department, a research department, and a human resources department.

Physical locations

Relecloud has one office in Seattle. The company has customers throughout North America and Europe.

Relecloud has most of its computing infrastructure located in a datacenter in Los Angeles. The company also has a few servers that are configured as Microsoft Azure virtual machines.

Existing Environment

Active Directory

The network contains an Active Directory forest named relecloud.com. The forest contains two domains named relecloud.com and relecloud.local. The functional level of the domain is Windows Server 2008 R2.

Each office contains five domain controllers. Each office is configured as an Active Directory site.

Current business model

Relecloud maintains a single Microsoft SQL Server database named Database1. Database1 hosts all of the data for the customers of Relecloud. Database1 has a total of 600 GB of data.

Relecloud provides data for its customers by using several methods of delivery, such as comma-separated value (CSV) files and fixed-width text files. All of the files are delivered by using File Transfer Protocol (FTP) or Secure (SFTP).

The SQL Server environment contains two on-premises servers and one server in the Azure region of US West. The Azure infrastructure contains an ExpressRoute connection.

The network contains the SQL servers configured as shown in the following table:

Server name	Location	CPU count	RAM
ServerDB1	On-premise	16	512GB
ServerDB2	On-premises	16	512GB
AzDB01	US West region	16	112 GB

ServerDB1 and ServerDB2 are members of an availability group.

Problem statements

Relecloud identifies several issues with its current data storage model. Multiple customers report that they have issues loading files into relational databases. Additionally, Relecloud administrators often fail to communicate with the FTP servers of the Relecloud customers.

Currently, the memory and the CPU of ServerDB1 and ServerDB2 are nearly fully utilized.

Requirements

Business goals

Relecloud plans to take a leadership role in its industry and modernize its delivery method. The company also plans to offer its customers easier access to its data.

Planned changes

Relecloud plans to move all of its customers to a delivery model that uses Azure SQL Database and plans to maintain a single database for each customer.

Relecloud plans to automate the process of creating the required infrastructure for each new customer, including the process of loading data to the customer databases.

Relecloud plans to onboard 30 new customer accounts during the next year. The data for the new customers is approximately 10 GB and has a static size. The new customers require a rapid response time to database queries and each customer has five to 10 concurrent logins.

Security requirements

Relecloud identifies the following requirements:

- Customers must be able to read data only.
- Customers must be able to create custom views.
- Customers must be able to use two-factor authentication when accessing Azure SQL databases.

Technical requirements

Relecloud identifies the following technical requirements:

- Database administrators must be notified when the CPU pressure of AzDB01 exceeds 80 percent.
- If ServerDB0 and ServerDB2 fail, AzDB01 must be able to respond to customer requests for data.

Auditing requirements

Relecloud identifies the following requirements:

- All of the requests from the stored procedures that are run against the Azure SQL databases must be logged.
- All auditing information must be viewable in Microsoft Excel in real-time.

A.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

DRAG DROP



You plan to prepare the environment for the planned new customers.

You need to deploy a base copy of DB1 for each customer in Azure.

Which cmdlets should you use in the Azure PowerShell script? To answer, drag the appropriate values to the correct targets. Each value 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.

Select and Place:

Answer Area

Get-AzureRmRecoveryServices
Backup

Get-AzureRmSqlDatabase

New-AzureRmSqlDatabase

Restore-AzureRmSqlDatabase

Set-AzureRmSiteRecovery
Protection

\$CustomerData = =ResourceGroupName "RG1"

-ServerName "server1" =DatabaseName "DG1"

-fromPointInTimeBackup -PointInTime UTCDateTime

-ResourceGroupName \$CustomerData.ResourceGroupName -ServerName
\$CustomerData.ServerName

-TargetDatabaseName "CustomerDatabase" -ResourceId \$CustomerData.ResourceId

-ElasticPoolName "EP1"

Correct Answer:

Get-AzureRmRecoveryServices
Backup

New-AzureRmSqlDatabase

Set-AzureRmSiteRecovery
Protection

Answer Area

```
$CustomerData = Get-AzureRmSqlDatabase -ResourceGroupName "RG1"
-ServerName "server1" -DatabaseName "DG1"

Restore-AzureRmSqlDatabase -fromPointInTimeBackup -PointInTime UTCDateTime
-ResourceGroupName $CustomerData.ResourceGroupName -ServerName
$CustomerData.ServerName
-TargetDatabaseName "CustomerDatabase" -ResourceID $CustomerData.ResourceID
-ElasticPoolName "EP1"
```

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Get-AzureRmSqlDatabase

The Get-AzureRmSqlDatabase cmdlet gets one or more Azure SQL databases from an Azure SQL Database Server.

Partial syntax:

Get-AzureRmSqlDatabase

[[[-DatabaseName] <String>]

[-ServerName] <String>

[-ResourceGroupName] <String>

Box 2: Restore-AzureRmSqlDatabase

The Restore-AzureRmSqlDatabase cmdlet restores a SQL database from a geo-redundant backup, a backup of a deleted database, a long term retention backup, or a point in time in a live database. The restored database is created as a new database.

Partial syntax:

Restore-AzureRmSqlDatabase

```
[ -FromLongTermRetentionBackup ]  
-ResourceId <String>  
-ServerName <String>  
-TargetDatabaseName <String>  
[ -Edition <String> ]  
[ -ServiceObjectiveName <String> ]  
[ -ElasticPoolName <String> ]  
[ -AsJob ]  
[ -LicenseType <String> ]  
[ -ResourceGroupName ] <String>  
[ -DefaultProfile <IAzureContextContainer> ]  
[ <CommonParameters> ]
```

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/get-azurermsqldatabase?view=azurermps-6.7.0>

<https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/restore-azurermsqldatabase?view=azurermps-6.7.0>

QUESTION 3

HOTSPOT

Users at Relecloud report show performance when viewing reports from AzDB01.

You need to identify the cause of the performance issue as quickly as possible.

Which dynamic management view and Performance Monitor counter should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Dynamic management view:

	▼
Sys.dm_io_cluster_shared_drives	
Sys.dm_io_virtual_file_stats	
Sys.dm_os_schedulers	

Performance Monitor counter:

	▼
Page Life Expectancy	
Page Reads/sec	
Processor Queue Length	

Correct Answer:

Answer Area

Dynamic management view:

	▼
Sys.dm_io_cluster_shared_drives	
Sys.dm_io_virtual_file_stats	
Sys.dm_os_schedulers	

Performance Monitor counter:

	▼
Page Life Expectancy	
Page Reads/sec	
Processor Queue Length	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Sys.dm_io_virtual_file_stats

Sys.dm_io_virtual_file_stats returns I/O statistics for data and log files.

Box 2:

Page reads/sec indicates the number of physical database page reads that are issued per second. This statistic displays the total number of physical page reads across all databases. Because physical I/O is expensive, you may be able to minimize the cost, either by using a larger data cache, intelligent indexes, and more

efficient queries, or by changing the database design.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-io-virtual-file-stats-transact-sql?view=sql-server-2017>

<https://docs.microsoft.com/en-us/sql/relational-databases/performance-monitor/monitor-cpu-usage?view=sql-server-2017>

QUESTION 4

You need to design a security strategy to support the planned changes for the customer databases. The solution must meet the security requirements.

Which two security roles should you grant to the customers? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. db_datawriter
- B. dbcreator
- C. db_owner
- D. db_ddladmin
- E. db_datareader

Correct Answer: DE

Section: (none)

Explanation



Explanation/Reference:

From Scenario: Security requirements

Relecloud identifies the following requirements:

- Customers must be able to read data only. (db_datareader)
- Customers must be able to create custom views. (db_ddladmin)
- Customers must be able to use two-factor authentication when accessing Azure SQL databases.

QUESTION 5

You need to implement a disaster recovery solution that meets the technical requirements. The solution must NOT require customers to modify their connection strings.

Which disaster recovery solution should you implement?

- A. Transactional Replication to Azure
- B. AlwaysOn Availability Groups
- C. Transaction Log Shipping

D. Failover cluster instances

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation

Scenario: Technical requirements

Relecloud identifies the following technical requirements:

- If ServerDB0 and ServerDB2 fail, AzDB01 must be able to respond to customer requests for data.

The Always On availability groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012 (11.x), Always On availability groups maximizes the availability of a set of user databases for an enterprise. An availability group supports a failover environment for a discrete set of user databases, known as availability databases, that fail over together. An availability group supports a set of read-write primary databases and one to eight sets of corresponding secondary databases.

Reference: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/always-on-availability-groups-sql-server?view=sql-server-2017>

QUESTION 6

You need to recommend a database strategy to support the planned new customers. The solution must minimize costs.

What should you recommend?

- A. A Basic database pool that has at least 1.200 eDTUs
- B. A Premium database pool that has at least 1.000 eDTUs
- C. 30 Basic databases
- D. 30 Premium P14 databases

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Scenario: Planned changes

Relecloud plans to move all of its customers to a delivery model that uses Azure SQL Database and plans to maintain a single database for each customer.

Relecloud plans to automate the process of creating the required infrastructure for each new customer, including the process of loading data to the customer databases.

Relecloud plans to onboard 30 new customer accounts during the next year. The data for the new customers is approximately 10 GB and has a static size. The new customers require a rapid response time to database queries and each customer has five to 10 concurrent logins.

Incorrect Answers:

B: A premium database pool would be more costly.

Reference: <https://azure.microsoft.com/en-us/pricing/calculator/?service=sql-database>

QUESTION 7

DRAG DROP

You need to implement a solution to meet the technical requirements for the database administrators.

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

Select and Place:



Actions

Enable diagnostics for the resource group that contains AzDB01.

Configure email notifications.

Configure Policy-Based Management.

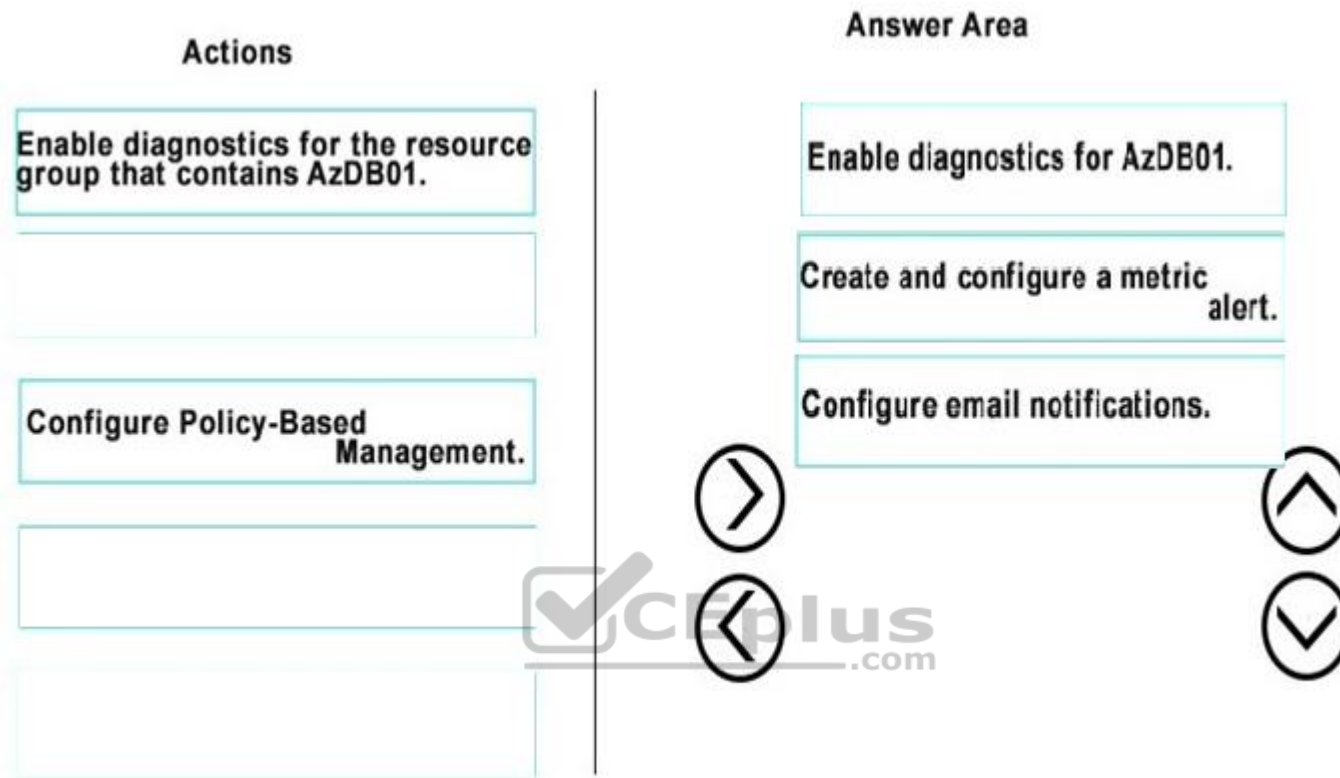
Enable diagnostics for AzDB01.

Create and configure a metric alert.

Answer Area



Correct Answer:



Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1:

Metrics and diagnostics logging is not enabled by default.

When you enable metrics and diagnostics logging, you need to specify the Azure resource where selected data is collected. Options available include: Log Analytics, Event Hubs, Storage

Example: To enable metrics and diagnostics logs collection in the portal, go to your SQL Database or elastic pool page, and then select Diagnostics settings. Etc.

Box 2:

You can receive an alert based on monitoring metrics for, or events on, your Azure services.

Box 3:

You can configure an alert to do the following when it triggers:

send email notifications to the service administrator and co-administrators

send email to additional emails that you specify.

call a webhook

Scenario:

Technical requirements

Relecloud identifies the following technical requirements:

- Database administrators must be notified when the CPU pressure of AzDB01 exceeds 80 percent.

Reference:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-insights-alerts-portal>

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-metrics-diag-logging>

QUESTION 8

Relecloud recently implemented the planned changes.

You plan to add a new customer named Fabrikam.Inc.

You need to ensure that the administrators at Fabrikam can access their data after the data loading process.

You collect the IP addresses used by the Fabrikam administrators.

Which Transact-SQL command or Azure PowerShell cmdlet should you use next?

- A. sp_set_firewall_rule
- B. Set-AzureRmSqlServerFirewallRule
- C. sp_set_database_firewall_rule
- D. New-AzureRmSqlServerFirewallRule

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:



Explanation:

The New-AzureRmSqlServerFirewallRule cmdlet creates a firewall rule for the specified Azure SQL Database server.

Database-level firewall rules enable clients to access certain (secure) databases within the same logical server. You can create these rules for each database (including the master database) and they are stored in the individual databases.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azurermsql/new-azurermsqlserverfirewallrule?view=azurermps-6.7.0>

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-firewall-configure>



Tailspin Toys

QUESTION 1

Overview

Tailspin Toys is a worldwide manufacturing company that operates on several continents.

The company has datacenters in London and New York.

Existing Environment

Technical Problems

Users in London office and New York office report quick response times when they connect to the datacenters. Users in all other office report slow response times.

Application Environment

Tailspin Toys stores sensitive customer data.

The company uses several applications. An application named App1 uses Microsoft SQL Server as a back-end database. An application named App2 uses MySQL as a back-end database.

Customer audits of documents rarely occur. When an audit occurs, it is scheduled between the customer and the auditing team of Tailspin Toys.

Requirements

Planned Changes

Tailspin Toys plans to migrate several workloads to Microsoft Azure. As part of the planned migration, the company requires that the lowest costs for monthly Azure services and for SQL Server license plans be implemented, whenever possible.

Technical Requirements

The customer-sensitive data stored in Azure must be encrypted when at rest.

The on-premises environment of Tailspin Toys has several jobs that support business processes scheduled by using the SQL Server Agent. After the migration to Azure, the business process must be scheduled to execute nightly at midnight Eastern Time (UTC-5). The implementation of the business processes must minimize administrative effort.

Networking requirements

All Internet traffic initiated from the servers in Azure must be routed through the local network firewalls and content filters.

Applications

App1

App1 is developed and maintained by the developers at Tailspin Toys. The company plans to make architectural changes to the application to achieve the best possible performance and scale.

The database for App1 is named DB1 and is 100-TB. 99.8 percent of the data in DB1 is large object binary (LOB) data. After the data is uploaded to DB1, the data is accessed frequently for 10 days; then the data is accessed rarely for 20 days; and then the data is never accessed, unless a customer audit occurs.

Users in multiple Tailspin Toys offices worldwide must use App1 to access the relational records stored in DB1.

App1 has a Service Level Agreement (SLA) that requires an uptime of 99.95 percent.

After the planned migration to Azure, the finance department plans to replicate a subset of DB1 to an on-premises database for reporting purposes.

App2

App2 is developed by using third-party open source code that is updated frequently. The database for App2 is named DB2. App2 only supports MySQL. None of the developers at Tailspin Toys are skilled in using MySQL. DB2 is supported by a team of external developers.

DB2 is 20 GB and is expected to reach 50 GB during the next five years. DB2 has a storage performance requirement of 200 IOPs.

DB2 must be highly available.

App3

As part of the planned migration to Azure, Tailspin Toys is developing a new application named App3 that will provide business-to-business features. Some of the customers that will use App3 will have a single employee, while others will have tens of thousands of employees using the application.

Due to the different sizes of the workloads, Tailspin Toys is concerned that parameter sniffing will hinder the performance of App3.

The web servers for App3 will be in multiple Azure regions and will use Traffic manager.

Customers worldwide will use App3. For each customer, the database for App3 must be located near the customer's corporate office. The customer also must be able to choose additional regions in which to store read-only copies of their database for an additional cost.

Two hundred customers plan to use App3 as soon as it becomes available. The customers request that their data be made available to them for reporting purposes. Some customers plan to use reporting in Azure, while other customers plan to use their own on-premises solution for reporting.

App3 has an SLA that requires 99.95 percent uptime.

The developers of App3 require an empty development database and a test database that contains data to test the application. The databases must be refreshed regularly.

A.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

QUESTION 2

You need to recommend a solution to meet the network requirements.

What should you recommend?

- A. Implement forced tunneling.
- B. Create a network security group (NSG) rule.
- C. Create an Azure virtual network gateway.
- D. Implement split tunneling.

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Application Gateway operates as an application delivery controller and offers SSL termination, cookie-based session affinity, round-robin load distribution, content-based routing, ability to host multiple websites and security enhancements.

Scenario: Networking requirements

All Internet traffic initiated from the servers in Azure must be routed through the local network firewalls and content filters.

Incorrect Answers:

B: If you need basic network level access control (based on IP address and the TCP or UDP protocols), you can use Network Security Groups (NSGs). An NSG is a basic, stateful, packet filtering firewall, and it enables you to control access based on a 5-tuple.

Reference: <https://docs.microsoft.com/en-us/azure/security/security-network-overview>

QUESTION 3

You need to recommend a solution to meet the reporting requirements of App3.

What should you include in the recommendation?

- A. SQL Server replication
- B. Database mirroring
- C. AlwaysOn Availability Groups
- D. Azure Data Sync

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

SQL Data Sync is a service built on Azure SQL Database that lets you synchronize the data you select bi-directionally across multiple SQL databases and SQL Server instances.

Scenario: Two hundred customers plan to use App3 as soon as it becomes available. The customers request that their data be made available to them for reporting purposes. Some customers plan to use reporting in Azure, while other customers plan to use their own on-premises solution for reporting.

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

QUESTION 4

HOTSPOT

You need to provision the database for DB2 in Azure.

How should you provision the database? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Service:

	▼
An Azure virtual machine that has MySQL installed	
Azure Cosmos DB	
Azure Data Lake	
Azure Database for MySQL	
Azure Database for PostgreSQL	
Azure SQL Database	

Tier:

	▼
Basic	
Premium	
Standard	

Correct Answer:

Answer Area

Service:

	▼
An Azure virtual machine that has MySQL installed	
Azure Cosmos DB	
Azure Data Lake	
Azure Database for MySQL	
Azure Database for PostgreSQL	
Azure SQL Database	

Tier:

	▼
Basic	
Premium	
Standard	

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Database for MySQL

Azure Database for MySQL provides fully managed, enterprise-ready community MySQL database as a service.

Box 2:

	Basic	General Purpose	Memory Optimized
Storage type	Azure Standard Storage	Azure Premium Storage	Azure Premium Storage
Storage size	5 GB to 1 TB	5 GB to 4 TB	5 GB to 4 TB
Storage increment size	1 GB	1 GB	1 GB
IOPS	Variable	3 IOPS/GB Min 100 IOPS Max 6000 IOPS	3 IOPS/GB Min 100 IOPS Max 6000 IOPS

Incorrect Answers:

Not Azure Database for MySQL

You can create an Azure Database for MySQL server in one of three different pricing tiers: Basic, General Purpose, and Memory Optimized.

Scenario: App2 is developed by using third-party open source code that is updated frequently. The database for App2 is named DB2. App2 only supports MySQL. None of the developers at Tailspin Toys are skilled in using MySQL. DB2 is supported by a team of external developers. DB2 is 20 GB and is expected to reach 50 GB during the next five years. DB2 has a storage performance requirement of 200 IOPS. DB2 must be highly available.

Reference:

<https://docs.microsoft.com/en-us/azure/mysql/concepts-pricing-tiers>

<https://azure.microsoft.com/en-us/services/mysql/>

QUESTION 5

DRAG DROP

You need to migrate DB2 to Azure.


Which utility should you use for each task? To answer, drag the appropriate utilities to the correct targets. Each utility 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.

Select and Place:

Utilities	Answer Area
azcopy	Export from on-premises DB2: <i>Utility</i>
bcp	Import to DB2 in Azure: <i>Utility</i>
mysql	
mysqldump	



Correct Answer:

Utilities	Answer Area
<input type="text"/>	Export from on-premises DB2: <input type="text" value="mysqldump"/>
<input type="text" value="bcp"/>	Import to DB2 in Azure: <input type="text" value="azcopy"/>
<input type="text" value="mysql"/>	
<input type="text"/>	

Section: (none)
Explanation

Explanation/Reference:
Explanation:

Box 1: mysqldump

Use common utilities and tools such as MySQL Workbench, mysqldump, Toad, or Navicat to remotely connect and restore data into Azure Database for MySQL.

Box 2: azcopy

AzCopy is a command-line tool for copying data to or from Azure Blob storage, Azure Files, and Azure Table storage, by using simple commands. The commands are designed for optimal performance. You can copy data between a file system and a storage account, or between storage accounts.

Explanation: App2 is developed by using third-party open source code that is updated frequently. The database for App2 is named DB2. App2 only supports MySQL.

Reference:

<https://docs.microsoft.com/en-us/azure/mysql/concepts-migrate-import-export>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-migrate-on-premises-data?tabs=windows>

QUESTION 6

What should you use to schedule the nightly jobs to meet the technical requirements?

- A. Migrate the SQL Server Agent jobs.
- B. Create an Azure Scheduler task.
- C. Create an Azure runbook.
- D. Implement an Elastic Database job.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Azure Automation can play the role of SQL Server Agent for Azure SQL DB.

Built into the Azure platform is Azure Automation and to get started, you have to create an automation account. You'll need to provide a name for the account, select your subscription, resource group, location, and determine if you want to create an Azure Run As account.

Once you create your account, you can then start creating runbooks. You can do just about anything with the runbooks. There are numerous existing run books that you can browse through and modify for your own use, including provisioning, monitoring, life cycle management, and more.

Scenario: After the migration to Azure, the business process must be scheduled to execute nightly at midnight Eastern Time (UTC-5). The implementation of the business processes must minimize administrative effort.

Incorrect Answers:

A: SQL Server agent is not supported in Azure.

Reference: <https://sqlperformance.com/2017/06/azure/automation-methods>

QUESTION 7

You need to prepare the Azure environment for the planned migration of DB1. The solution must meet the requirements of the finance department.

Which two tasks should you perform in Azure? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Configure a Filtered replication publication
- B. Create an Azure virtual machine that runs SQL Server
- C. Configure multiple active replicas.
- D. Create an Azure SQL database.

Correct Answer: AD

Section: (none)

Explanation**Explanation/Reference:**

References: <https://docs.microsoft.com/en-us/sql/relational-databases/replication/publish/filter-published-data?view=sql-server-2017>

QUESTION 8

You need to implement a solution to meet the technical requirements for the encryption of the data in App1.

Which two features should you enable? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Transparent Data Encryption (TDE)
- B. Storage Account Encryption
- C. Dynamic Data Masking
- D. Azure Key Vault

Correct Answer: AD

Section: (none)

Explanation**Explanation/Reference:**

An application named App1 uses Microsoft SQL Server as a back-end database. The customer-sensitive data stored in Azure must be encrypted when at rest.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/extensible-key-management-using-azure-key-vault-sql-server?view=sql-server-2017>

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/transparent-data-encryption?view=sql-server-2017>

QUESTION 9

Which database platform should App3 use?

- A. Azure Database for MySQL
- B. Azure SQL Database
- C. Azure Database for PostgreSQL
- D. an Azure virtual machine that has SQL Server installed

Correct Answer: A

Section: (none)

Explanation**Explanation/Reference:**

QUESTION 10

DRAG DROP

You need to identify the appropriate data architecture for App1. The solution must support the planned changes.

Which storage platforms should you identify for App1? To answer, drag the appropriate storage platforms to the correct data types. Each storage platform 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.

Select and Place:

Storage Platforms

An Azure virtual machine that has SQL Server installed

Azure Cosmos DB

Azure Data Lake

Azure Database for MySQL

Azure Database for PostgreSQL

Azure SQL Database

Answer Area

Relational data:

LOB data:

Storage platform

Storage platform

Correct Answer:

Storage Platforms

An Azure virtual machine that has SQL Server installed

Azure Cosmos DB

Azure Database for MySQL

Azure Database for PostgreSQL

Answer Area

Relational data: Azure SQL Database

LOB data: Azure Data Lake

Section: (none)
Explanation

Explanation/Reference:

Mixed Questions

QUESTION 1

You have Microsoft SQL Server 2014 installed on a Microsoft Azure virtual machine.

One of the databases on the virtual machine supports a highly active Online Transaction Processing (OLTP) application. Users report abnormally long wait times when they submit data in the application.

Which two tools can you use to identify the longest running queries? Each correct answer presents a complete solution.

- A. the Job Activity Monitor
- B. Database Engine Tuning Advisor
- C. dynamic management views
- D. SQL Server Extended Events
- E. SQL metrics in Azure Diagnostics for the virtual machine

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

C: Microsoft Azure SQL Database enables a subset of dynamic management views to diagnose performance problems, which might be caused by blocked or long-running queries, resource bottlenecks, poor query plans, and so on. You can detect such common performance problems by using dynamic management views.

D: The SQL Profiler has been replaced by SQL Server Extended Events.

Extended Events works via Event Tracing (ETW). This has been the common way for all Microsoft related technologies to expose diagnostic data. ETW provides much more flexibility.

Reference:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-monitoring-with-dmvs> <https://stackify.com/performance-tuning-in-sql-server-find-slow-queries>

QUESTION 2

You have an application that uses a Microsoft SQL Server database on a Microsoft Azure virtual machine.

The application experiences performance issues, which you suspect are related to the connection pooling. The issues are prevalent only when there are more than 150 concurrent connections.

You need to identify whether the performance issues are related to the connection pooling.

Which dynamic management view should you query?

- A. sys.dm_os_memory_pools
- B. sys.dm_exec_requests
- C. sys.dm_exec_connections
- D. sys.dm_exec_sessions

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

References: https://blogs.msdn.microsoft.com/sql_pfe_blog/2013/10/08/connection-pooling-for-the-sql-server-dba/

QUESTION 3

A customer plans to monitor the performance of a Microsoft Azure SQL database.

You need to explain to the customer which metrics are used to calculate the Database Throughput Unit (DTU) percentage. Which three metrics should you identify? Each correct answer presents part of the solution.

- A. Log IO percentage
- B. Data IO percentage
- C. CPU percentage
- D. Blocked by firewall
- E. Database size percentage
- F. Total database size

Correct Answer: ABC

Section: (none)

Explanation

Explanation/Reference:

References:

<http://dtucalculator.azurewebsites.net/>

QUESTION 4

You have a Microsoft Azure SQL database that is used for reporting.

You discover that some reports complete more quickly than other reports, even though the reports retrieve approximately the same amount of data and use queries that have the same table structure.

You need to reduce the amount of time it takes to complete the reports.

What should you use?

- A. Index Tuning Wizard
- B. Database Engine Tuning Advisor
- C. Azure Throughput Analyzer



D. Azure SQL Database Index Advisor

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 5

You plan to implement a Microsoft Azure SQL database.

You need to create and manage the new database on a new server.

Which three cmdlets should you use? Each correct answer presents part of the solution.

- A. New-AzureSqlDatabaseServer
- B. New-AzureSqlDatabaseServerFirewallRule
- C. New-AzureSqlDatabaseServerContext
- D. New-AzureVM
- E. New-AzureSqlDatabase

Correct Answer: ACE

Section: (none)

Explanation

Explanation/Reference:

References: <https://msdn.microsoft.com/en-us/library/dn546722.aspx>

QUESTION 6

You have a Microsoft Azure SQL database. The database is hosted in the West US region and uses the Premium service tier. Users of the database are located in Los Angeles, New York, and Singapore.

The users in Singapore report that when they run reports against the database, the reports take a long time to complete. The reports contain thousands of rows. You need to recommend a solution to resolve the performance issue. The solution must maintain the performance for the other users. What should you include in the recommendation?

- A. Move the Azure SQL database from the West US region to the East Asia region.
- B. Implement Azure ExpressRoute for the subscription.
- C. Configure a readable geo-replica in the East Asia region.
- D. Configure pagination for the report.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

QUESTION 7

You have an on-premises Microsoft SQL Server 2014 database in an AlwaysOn availability group.

You are planning a backup solution for the database. Backups will run on a secondary replica.

You need to create a backup procedure for the database. The solution must ensure that the backups are stored on-premises and in Microsoft Azure Blob storage.

Which two actions should you perform? Each correct answer presents part of the solution.

- A. Run the BACKUP DATABASE statement and use the TO DISK clause.
- B. Run the CREATE DATABASE statement and use the AS COPY OF clause.
- C. Run the BACKUP DATABASE statement and use the TO URL clause.
- D. Run the BACKUP DATABASE statement and use the MIRROR clause.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

<https://msdn.microsoft.com/en-us/library/dn435916.aspx>

QUESTION 8

You have a hybrid Microsoft SQL Server environment that has multiple servers and services that run in both Microsoft Azure and on-premises.

If the network fails, you need to ensure that users can authenticate and connect to the resources available in Azure.

Which two server roles should you deploy to Azure? Each correct answer presents part of the solution.

- A. Active Directory Federation Services
- B. Active Directory Certificate Services
- C. DHCP Server
- D. Active Directory Domain Services
- E. DNS Server

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 9

You have a Microsoft Azure SQL data warehouse.

During peak usage, you discover that the data warehouse fails to meet performance expectations.

You need to provide additional resources to the data warehouse.

What should you do?

- A. Create an additional storage pool.
- B. Create an elastic database pool.
- C. Increase the number of Database Throughput Units (DTUs).
- D. Increase the number of Data Warehouse Units (DWUs).

Correct Answer: D

Section: (none)

Explanation



Explanation/Reference:

References: <https://azure.microsoft.com/en-gb/documentation/articles/sql-data-warehouse-manage-compute-overview/>

QUESTION 10

You have an organizational data mart that contains 3.5 TB of uncompressed data.

The size of the data is expected to grow 2 TB annually. Reporting is performed only during business hours.

You plan to move the data to the cloud.

You need to design a cloud architecture for the data mart. The solution must minimize costs and administrative overhead.

What is the best design to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. a Microsoft Azure SQL data warehouse
- B. a Microsoft SQL Server virtual machine
- C. a Microsoft Azure HDInsight cluster
- D. a Microsoft Azure SQL database

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

References: <https://azure.microsoft.com/en-gb/documentation/articles/sql-data-warehouse-overview-what-is/>

QUESTION 11

Your company has two offices. The offices are located in London and New York.

The London office has a public IP of 131.107.1.1 and a private IP subnet of 192.168.4.0/24. The New York office has a public IP of 131.107.15.1 and a private IP subnet of 192.168.8.0/24.

You have a Microsoft SQL Server database that is hosted on a Microsoft Azure virtual machine. Currently, only the users in the London office use the database. You need to create a new firewall rule to provide the users in the New York office with access to the database. The solution must protect the database from being accessed by users on the public Internet.

What should you do?

- A. Modify the endpoint from the Azure portal.
- B. Use the `sp_set_database_firewall_rule` stored procedure.
- C. Disable the Windows Firewall on the Azure virtual machine.
- D. Run the `New-AzureSqlDatabaseServerFirewallRuleAzure` PowerShell cmdlet.
- E. Use the Set Firewall Rule REST API.

Correct Answer: A

Section: (none)

Explanation

**Explanation/Reference:**

This Microsoft SQL Server database is hosted on a Microsoft Azure virtual machine. Therefore, remote access will be provided via the endpoint rather than a database level firewall rule.

QUESTION 12

You have a Microsoft SQL Server instance on a Microsoft Azure virtual machine.

The members of an Active Directory group named HelpDesk can log in to the SQL Server instance.

You need to ensure that the members of HelpDesk can query dynamic management views and gather performance metrics from the SQL Server instance.

Which three actions should you perform? Each correct answer presents part of the solution.

- A. Add HelpDesk to the `db_owner` role for all of the databases.
- B. Add HelpDesk to the new role.
- C. Grant `VIEW ANY DATABASE` to the new role.
- D. Grant `VIEW SERVER STATE` to the new role.
- E. Create a database role.

F. Create a server role.

Correct Answer: BDF

Section: (none)

Explanation

Explanation/Reference:

References: <https://msdn.microsoft.com/en-us/library/ms188754.aspx>

QUESTION 13

You have a Microsoft Azure SQL database.

You need to ensure that you can use the Always Encrypted feature for the database.

Which two components should you configure? Each correct answer presents part of the solution.

- A. the Column Master Key
- B. the database master key
- C. the service master key
- D. transparent data encryption (TDE)
- E. the Column Encryption Key

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

References: <https://azure.microsoft.com/en-gb/documentation/articles/sql-database-always-encrypted/>

QUESTION 14

Your company has two offices located in London and New York.

The London office has a public IP of 131.107.1.1 and a private IP subnet of 192.168.4.0/24. The New York office has a public IP of 131.107.15.1 and a private IP subnet of 192.168.8.0/24.

You have a Microsoft Azure SQL database. Currently, only the users in the London office use the database.

You plan to provide the users in the New York office with access to the database.

You need to ensure that the New York office users can access the database.

Which command should you execute?

- A. EXECUTE sp_set_database_firewall_rule N'NewYork','131.107.15.1','255.255.255.255'
- B. EXECUTE sp_set_database_firewall_rule N'NewYork, '192.168.8.0', '192.168.8.255';
- C. EXECUTE sp_set_database_firewall_rule N'NewYork,'192.168.8.0', '255.255.255.0';

D. EXECUTE sp_set_database_firewall_rule N'NewYork','131.107.15.1','131.107.15.1';

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

References:

<http://azure.patrickkeisler.com/2015/11/configuring-the-azure-sql-database-firewall/> <https://msdn.microsoft.com/en-gb/library/dn270010.aspx>

QUESTION 15

You have a Microsoft SQL Server 2014 instance on a Microsoft Azure virtual machine.
You need to ensure that all SQL Server audits are written to the Windows Security log.
Which two actions should you perform? Each correct answer presents part of the solution.

- A. Configure the Audit policy change policy.
- B. Assign the Generate security audits user right.
- C. Configure the Audit object access policy.
- D. Assign the Log on as a service user right.
- E. Configure the Audit the access of global system objects policy.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

References: <https://msdn.microsoft.com/en-us/library/cc645889.aspx>

QUESTION 16

In this section you will see one or more sets of questions with the same scenario and problem. Each question-presents a unique solution to the problem, and you must determine whether the solution meets the stated goals. More than one solution might solve the problem. It is also possible that none of the solutions solve the problem.

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.

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 stayed 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 on the review screen.

You have a Microsoft Azure Database for MySQL server named SQLDB1.

Four database administrators manage SQLDB1.

You need to prevent any administrator from deleting the server that hosts SQLDB1.

Solution: You create a delete lock on the server resource in Azure.
Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

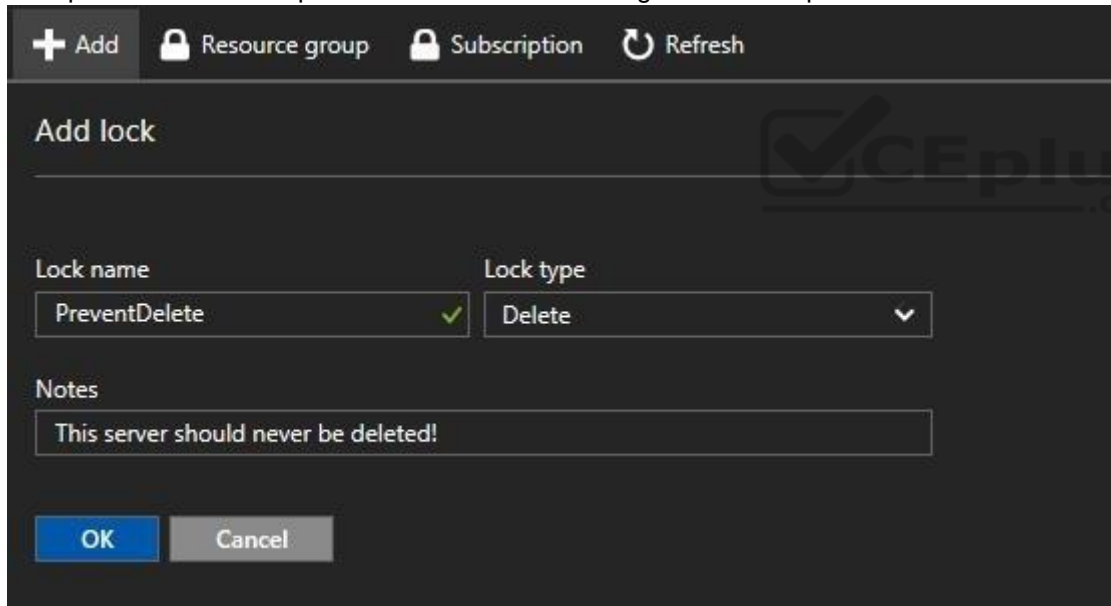
Section: (none)

Explanation

Explanation/Reference:

As an administrator, you may need to lock a subscription, resource group, or resource to prevent other users in your organization from accidentally deleting or modifying critical resources. You can set the lock level to CanNotDelete or ReadOnly. In the portal, the locks are called Delete and Read-only respectively. It can help prevent accidental deletion resources such as a SQL Server by adding a level of protection.

Example: A lock that will prevent the server from being deleted and provide a note that it should never be deleted.



The screenshot shows the 'Add lock' dialog in the Azure portal. The dialog has a title bar with '+ Add', 'Resource group', 'Subscription', and 'Refresh' buttons. The main content area is titled 'Add lock'. It contains a 'Lock name' field with the text 'PreventDelete' and a green checkmark icon to its right. To the right of the 'Lock name' field is a 'Lock type' dropdown menu with 'Delete' selected. Below these fields is a 'Notes' text area containing the text 'This server should never be deleted!'. At the bottom of the dialog are two buttons: 'OK' and 'Cancel'.

References: <https://blogs.msdn.microsoft.com/azuresqldbssupport/2017/06/19/protecting-deletions-of-azure-sql-resources/>

QUESTION 17

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 stayed 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 on the review screen.

You have a Microsoft Azure Database for MySQL server named SQLDB1.

Four database administrators manage SQLDB1.

You need to prevent any administrator from deleting the server that hosts SQLDB1.

Solution: You remove the Delete permission for SQLDB1 from the Azure portal.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Use the Delete lock instead.

References: <https://blogs.msdn.microsoft.com/azuresqldbssupport/2017/06/19/protecting-deletions-of-azure-sql-resources/>

QUESTION 18

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 stayed 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 on the review screen.

You have a Microsoft Azure Database for MySQL server named SQLDB1.

Four database administrators manage SQLDB1.

You need to prevent any administrator from deleting the server that hosts SQLDB1.

Solution: You create a delete lock on the database resource in Azure.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

The delete lock should on the server resource instead.

References: <https://blogs.msdn.microsoft.com/azuresqldbssupport/2017/06/19/protecting-deletions-of-azure-sql-resources/>

QUESTION 19

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 stayed 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 on the review screen. You plan to configure auditing for a Microsoft SQL Server instance that contains 17 databases. You need to audit all of the SELECT statements that are executed against all of the databases. The solution must ensure that all gathered information is written to the application log.
Solution: You create a Database Audit Specification in each database.
Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

An auditing policy can be defined for a specific database or as a default server policy. A server policy applies to all existing and newly created databases on the server.

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

QUESTION 20

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 stayed 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 on the review screen. You plan to configure auditing for a Microsoft SQL Server instance that contains 17 databases. You need to audit all of the SELECT statements that are executed against all of the databases. The solution must ensure that all gathered information is written to the application log.
Solution: You create a Database Audit Specification in the master database.
Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

An auditing policy can be defined for a specific database or as a default server policy. A server policy applies to all existing and newly created databases on the

server.

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

QUESTION 21

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 stayed 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 on the review screen.

You plan to configure auditing for a Microsoft SQL Server instance that contains 17 databases.

You need to audit all of the SELECT statements that are executed against all of the databases. The solution must ensure that all gathered information is written to the application log.

Solution: You create a Server Audit Specification in the master database.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

An auditing policy can be defined for a specific database or as a default server policy. A server policy applies to all existing and newly created databases on the server.

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

QUESTION 22

You have a Microsoft Azure SQL database named DB1 and a copy named DB1Test.

You use DB1Test to test the deployment of major updates to the database.

You discover that DB1Test contains outdated data.

You need to ensure that DB1Test contains recent data from DB1.

Which two statements should you execute? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. CREATE DATABASE DB1Test AS Snapshot OF DB1

B. CREATE DATABASE DB1Test AS COPY OF DB1

C. SELECT d.name AS DBName, d.state-desc
FROM sys.databases d
WHERE name= 'DB1Test'

D. DROP DATABASE DB1Test

E. ALTER DATABASE DB1TestSET ONLINE;

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

This command copies Database1 to a new database named Database2 on the same server. Depending on the size of your database, the copying operation might take some time to complete.

-- Execute on the master database.

-- Start copying.

CREATE DATABASE Database1_copy AS COPY OF Database1;

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

QUESTION 23

You manage a large Microsoft Azure SQL database. The database contains data that changes several thousand times each day.

You need to ensure that the performance of the database stays consistent when the data changes.

What should you do?

- A. Update the statistics of all the tables in the database.
- B. Clear the procedure cache of the database.
- C. Run a consistency check on the tables in the database.
- D. Enable Extended Events for the server that contains the database.



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Updating statistics ensures that queries compile with up-to-date statistics.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/update-statistics-transact-sql?view=sql-server-2017>

QUESTION 24

Your company has a Microsoft SQL Server database hosted on a Microsoft Azure virtual machine.

You create several reports that display data from the database.

Users report that the reports are slow to display data. You create an index.

You need to identify whether the performance of the reports significantly improves.

- A. SQL Server Profiler

- B. Query Store
- C. Database Engine Tuning Advisor
- D. Live Query Statistics

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

To view the live query execution plan, on the tools menu click the Live Query Statistics icon.

The live query plan displays the overall query progress and operator-level run-time execution statistics such as the number of rows produced, elapsed time, operator progress, etc. Because this data is available in real time without needing to wait for the query to complete, these execution statistics are extremely useful for debugging query performance issues.

Incorrect Answers:

A: SQL Server Profiler is being discontinued.

B: The SQL Server Query Store feature provides you with insight on query plan choice and performance. It simplifies performance troubleshooting by helping you quickly find performance differences caused by query plan changes. Query Store automatically captures a history of queries, plans, and runtime statistics, and retains these for your review. It separates data by time windows so you can see database usage patterns and understand when query plan changes happened on the server.

C: Database Engine Tuning Advisor examines how queries are processed in the databases you specify, and then recommends how you can improve query processing performance by modifying database structures such as indexes, indexed views, and partitioning.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/live-query-statistics?view=sql-server-2017>

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store?view=sql-server-2017> <https://docs.microsoft.com/en-us/sql/tools/dta/tutorial-database-engine-tuning-advisor?view=sql-server-2017>

QUESTION 25

Your company has a data warehouse that contains a database named DB1.

You plan to implement data encryption DB1.

You need to configure DB1 to meet the following requirements:

- Prevent the system administrators of App1 from reading the Personally Identifiable Information (PIII) stored in DB1.

- Ensure that all backups are stored in an on-premises local drive for three days.

- Ensure that all backups are stored in a remote location for 10 days. Encrypt all backups.

All solutions must support the ODBC Driver for Microsoft SQL Server 2012.

Which two solutions should you implement for DB1? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Encrypt within the application tier.
- B. Backup DB1 by using VSS snapshots.
- C. Implement the Always Encrypted feature.

- D. Backup DB1 to a disk and a URL by using MIRROR TO.
- E. Encrypt within DB1 by using EncryptByPassword.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

The backup command argument has the MIRROR TO option

Always Encrypted is a new data encryption technology in Azure SQL Database and SQL Server that helps protect sensitive data at rest on the server, during movement between client and server, and while the data is in use. Always Encrypted ensures that sensitive data never appears as plaintext inside the database system.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault>

<https://docs.microsoft.com/en-us/sql/relational-databases/backup-restore/sql-server-backup-to-url?view=sql-server-2017>

QUESTION 26

You plan to create a data warehouse in Microsoft Azure for development purposes.

The data warehouse must not exceed 10 terabytes (TB) in size. You must minimize costs for the environment.

You need to create the data warehouse.

Which Transact-SQL statement should you run?

- A. CREATE DATABASE CloudDataWarehouse (
 EDITION = 'premium',
 SERVICE_OBJECTIVE = 'P1'
)
- B. CREATE DATABASE CloudDataWarehouse (
 EDITION = 'datawarehouse',
 SERVICE_OBJECTIVE = 'DW200'
)
- C. CREATE DATABASE CloudDataWarehouse (
 EDITION = 'basic'
 SERVICE_OBJECTIVE = 'basic'
)
- D. CREATE DATABASE CloudDataWarehouse (
 EDITION = 'standard',
 SERVICE_OBJECTIVE = '53')

Correct Answer: B

Section: (none)

Explanation



Explanation/Reference:

EDITION specifies the service tier of the database. For SQL Data Warehouse use 'datawarehouse'.

References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-database-transact-sql?view=sql-server-2017&tabs=sql_dw

QUESTION 27

You have 10 Microsoft Azure SQL database.

One of the databases is named DB1.

You need to ensure that only a specific IP address can be used to access DB1. The solution must affect only DB1.

What should you use?

- A. the `sp_set_database_firewall_rule` statement
- B. the `New-AzureSqlDatabaseServerFirewallRule` cmdlet
- C. the Azure portal to add an endpoint
- D. the `Set-AzureSqlDatabaseServerFirewallRule` cmdlet

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Section: [none]

Explanation



If the IP address of the request is within one of the ranges specified in the database-level firewall rules, the connection is granted to the SQL Database that contains the rule.

The `New-AzureSqlDatabaseServerFirewallRule` cmdlet creates a firewall rule in the specified instance of Azure SQL Database Server in the current subscription.

Use the `StartIpAddress` and `EndIpAddress` parameters to specify a range of IP addresses that this rule allows to connect to the Azure SQL Database server.

Specify the `AllowAllAzureServices` parameter to create a rule that allows Azure connections to the server. The rule has starting and ending IP address values of 0.0.0.0. If you do not specify a firewall rule name, this cmdlet assigns the default name `AllowAllAzureServices`.

References: <https://docs.microsoft.com/en-us/powershell/module/servicemanagement/azure/new-azuresqldatabaseserverfirewallrule?view=azuresmps-4.0.0>

QUESTION 28

Your company has a Microsoft Azure MySQL database that is used by an Azure Web App. The logical server that hosts the database contains other databases.

You move the database to another Azure tenant.

You need to ensure that the Web App can access the database.

What should you do?

- A. Modify the SSL settings for the database server.

- B. Add an Azure virtual network gateway.
- C. Create a server-level firewall rule.
- D. Create a database-level firewall rule.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Azure Database for MySQL provides access security using a firewall to protect your data.

Create a firewall rule

On the MySQL server blade, under the Settings heading, click Connection Security to open the Connection Security blade for Azure Database for MySQL.

References: <https://docs.microsoft.com/en-us/azure/mysql/howto-connect-webapp>

QUESTION 29

Your environment contains a Microsoft Azure virtual machine that has Microsoft SQL Server Standard installed.

An application named App1 uses the data stored in SQL Server. The Service Level Agreement (SLA) for App1 states that the application can lose a maximum of 10 minutes of data and that the data must be available if two Azure regions fail.

You need to implement a solution that provides disaster recovery for App1.

In two other Azure regions, you deploy an Azure virtual machine that has SQL Server installed.

What should you implement next?

- A. Merge Replication
- B. AlwaysOn Failover Cluster Instances
- C. AlwaysOn Availability Groups
- D. Log Shipping

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Availability Groups example: Availability replicas running across multiple datacenters in Azure VMs for disaster recovery. This cross-region solution protects against complete site outage.

Note: Microsoft SQL Server AlwaysOn is a collection of high availability and disaster recovery features introduced from SQL Server 2012. AlwaysOn is used to minimize the Recovery Point Objective (RPO) and Recovery Time Objective (RTO) and maximized availability databases. SQL Server AlwaysOn branding encompasses both FCIs (Failover Cluster) and AGs (Availability Group) an enterprise-level alternative to database mirroring.

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-sql-high-availability-dr> <https://docs.microsoft.com/en-us/azure/>

availability-zones/az-overview

QUESTION 30

You have a failover cluster deployed to Microsoft Azure virtual machines.

You plan to deploy a new Microsoft SQL instance to the cluster. The instance will host one database that has 10 file groups.

You provision 10 disks for the instance.

You need to recommend a storage solution that always provides the fastest possible reads for the database.

What should you include in the recommendation?

- A. Dynamic disks that use striping
- B. One file group per disk
- C. A Storage space
- D. Dynamic disks that use mirroring

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Dynamic disks provide features that basic disks do not, such as the ability to create volumes that span multiple disks (spanned and striped volumes) and the ability to create fault-tolerant volumes (mirrored and RAID-5 volumes). Disk striping can speed up operations that retrieve data from disk storage.

References: <https://docs.microsoft.com/en-us/windows/desktop/fileio/basic-and-dynamic-disks>

QUESTION 31

You plan to deploy a data platform solution for 4 TB of data. The solution will be used to generate analytic reports during business hours only. All reports will be generated by using Microsoft Excel.

You need to recommend a data platform for the data. The solution must minimize costs and provide the highest read performance.

Which data platform should you recommend?

- A. Microsoft Azure SQL Data Warehouse
- B. Microsoft Azure Cosmos DB
- C. Microsoft Azure SQL Database
- D. Microsoft Azure Data Lake

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

SQL Data Warehouse is a cloud-based Enterprise Data Warehouse (EDW) that leverages Massively Parallel Processing (MPP) to quickly run complex queries across petabytes of data. Use SQL Data Warehouse as a key component of a big data solution. Import big data into SQL Data Warehouse with simple PolyBase T-SQL queries, and then use the power of MPP to run high-performance analytics.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-overview-what-is>

QUESTION 32

You have a Microsoft Azure Infrastructure as a Service (IaaS) environment that contains a 3-TB data warehouse.

You need to recommend a solution to improve the read performance of the database. The solution must maintain data integrity.

Which task should you include in the recommendation?

- A. Enable read caching.
- B. Move the transaction logs to the temporary drive.
- C. Enable read/write caching.
- D. Set the interleave size to 64 KB.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Best practices include enabling read caching on the disk(s) hosting the data files and TempDB data files.

References: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-sql-performance>

QUESTION 33

Your company registers the domain name of contoso.com.

You plan to deploy two web servers, two application servers, and two Microsoft SQL Server instances on Microsoft Azure virtual machines. All of the virtual machines will be located on the same Azure virtual network. The SQL Server instances will be in an AlwaysOn availability group.

You need to ensure that the web servers and the application servers can access the availability group listener by using a fully qualified domain name (FQDN) from the contoso.com domain.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Register a public DNS service on the Azure virtual network.
- B. Add a subnet to the Azure virtual network.
- C. Deploy a domain controller and join the virtual machines to the domain.
- D. Register a DNS server on the Azure virtual network.
- E. Use the Azure-provided name resolution service.
- F. Deploy a DNS server.

Correct Answer: BCD

Section: (none)

Explanation

Explanation/Reference:

B: You need a subnet for the VMs.

C: You need a domain for the availability group.

D: Azure assigns IP addresses and DNS server addresses to the VMs. By default, Azure will assign an IP address of an Azure DNS server to the VMs on the network. To be able to join the VMs to a domain, you need the VMs to use the domain controller for their DNS server. You need to register a DNS server on the Azure virtual network. The DNS server you will register is the domain controller. Registering the domain controller's IP address as the DNS server will configure Azure to assign the domain controller's IP address to the other VMs in the network which will enable you to join them to the domain.

QUESTION 34

You implement Microsoft Azure ExpressRoute.

You need to ensure that if the ExpressRoute connection is unavailable, you can manage the resources in Azure over a secure connection.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Create an Azure virtual network.
- B. Enable a point-to-site VPN on the gateway.
- C. Install the Azure VPN client.
- D. Create a certificate.
- E. Create a private IP address.



Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

You can configure a Site-to-Site VPN as a secure failover path for ExpressRoute.

Static route should be configured for your VPN gateway. If your local network is connected to both ExpressRoute and a Site-to-Site VPN, you must have a static route configured in your local network to route the Site-to-Site VPN connection to the public Internet.

Reference: <https://docs.microsoft.com/en-us/azure/expressroute/expressroute-howto-coexist-resource-manager>

QUESTION 35

You recently implemented a Microsoft Azure SQL data warehouse.

You discover that during peak hours, the data warehouse takes longer than expected to generate reports.

You need to reduce the amount of time it takes to generate the reports.

What should you do from the Azure portal?

- A. Increase the number of Database Throughput Units (DTUs).
- B. Move the data warehouse to a different resource group.
- C. Add storage to the data warehouse.
- D. Increase the number of Data Warehouse Units (DWUs).

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

With SQL Data Warehouse CPU, memory, and IO are bundled into units of compute scale called Data Warehouse Units (DWUs). A DWU represents an abstract, normalized measure of compute resources and performance. By changing your service level you alter the number of DWUs that are allocated to the system, which in turn adjusts the performance, and the cost, of your system.

Incorrect Answers:

A: Database Throughput Unit (DTU): DTUs provide a way to describe the relative capacity of a performance level of Basic, Standard, and Premium databases. DTUs are based on a blended measure of CPU, memory, reads, and writes. As DTUs increase, the power offered by the performance level increases.

Reference: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/what-is-a-data-warehouse-unit-dwu-cdwu>

QUESTION 36

You need to deploy a new Microsoft Azure SQL database that must meet the following requirements:

Will be available if a single Azure region becomes unavailable Will be able to fail over to another Azure region automatically
Which two configuration tasks should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Add the database to a failover group.
- B. Add the database to a database pool.
- C. Configure the database connection string to point to the endpoint of the read-write listener.
- D. Add the database to an AlwaysOn availability group that uses synchronous replication.
- E. Configure the database connection string to point to the endpoint of the primary replica.

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

C: An availability group requires a load balancer when the SQL Server instances are on Azure virtual machines. The load balancer stores the IP address for the availability group listener. If an availability group spans multiple regions, each region needs a load balancer.

D: To test listener connectivity to the remote region, you can fail over the replica to the remote region. While the replica is asynchronous, failover is vulnerable to potential data loss. To fail over without data loss, change the availability mode to synchronous and set the failover mode to automatic.

Reference: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sql/virtual-machines-windows-portal-sql-availability-group-dr>

QUESTION 37

You have a Microsoft Azure SQL database named DB1.
You need to monitor DB1 to identify all regressed queries.
Which command should you use?

- A. `select*from sys.dm_exec_query_stats`
- B. `ALTER DATABASE DB1 SET QUERY_STORE (OPERATION_MODE=READ_WRITE)`
- C. `select*from sys.dm_exec_query_stats`
`cross apply sys.dm_exec_query_plan(plan_handle)`
- D. `ALTER DATABASE DB1 SET QUERY_STORE (QUERY_CAPTURE_MODE= ALL)GO`

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Queries with multiple plans? These queries are especially interesting because they are candidates for regressions due to plan choice change. The following query identifies these queries along with all plans:

WITH Query_MultPlans

AS

```
(  
SELECT COUNT(*) AS cnt, q.query_id  
FROM sys.query_store_query_text AS qt  
JOIN sys.query_store_query AS q  
    ON qt.query_text_id = q.query_text_id  
JOIN sys.query_store_plan AS p  
    ON p.query_id = q.query_id  
GROUP BY q.query_id  
HAVING COUNT(distinct plan_id) > 1 )
```

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store?view=sql-server-2017#Regressed>

<https://blogs.msdn.microsoft.com/sqlserverstorageengine/2017/04/24/how-to-find-query-plan-choice-regressions-with-sql-server-2017-ctp2/>

QUESTION 38

You need to create a Microsoft Azure Resource Manager template.
Which file format should you use?

- A. XML

- B. text
- C. JSON
- D. binary

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

An Azure Resource Manager template presents the different sections of a template and the properties that are available in those sections. The template consists of JSON and expressions that you can use to construct values for your deployment.

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

QUESTION 39

You have Microsoft SQL Server installed on a Microsoft Azure virtual machine.

One of the databases on the virtual machine supports a highly active Online Transaction Processing (OLTP) application. Users report abnormally long wait times when they submit data in the application.

Which two tools can you use to identify the longest running queries? Each correct answer presents a complete solution.

- A. Dynamic management views
- B. SQL Server Extended Events
- C. Database Engine Tuning Advisor
- D. the Job Activity Monitor
- E. SQL metrics in Azure Diagnostics for the virtual machine

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

A: Microsoft Azure SQL Database enables a subset of dynamic management views to diagnose performance problems, which might be caused by blocked or long-running queries, resource bottlenecks, poor query plans, and so on. You can detect such common performance problems by using dynamic management views.

B: The SQL Profiler has been replaced by SQL Server Extended Events.

Extended Events works via Event Tracing (ETW). This has been the common way for all Microsoft related technologies to expose diagnostic data. ETW provides much more flexibility.

Reference: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-monitoring-with-dmvs> <https://stackify.com/performance-tuning-in-sql-server-find-slow-queries>

QUESTION 40

You have a Microsoft Azure SQL database named DB1.

You plan to monitor the resource utilization of DB1.

You need to review the following resource utilizations:

Average compute Average I/O

Which data management view should you use?

- A. sys.dm_db_resources_stats
- B. sys.elastic_pool_resource_stats
- C. sys.dm_exec_query_stats
- D. sys.resource_stats

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Section: [none]

Explanation

sys.dm_db_resource_stats (Azure SQL Database) returns CPU, I/O, and memory consumption for an Azure SQL Database database. One row exists for every 15 seconds, even if there is no activity in the database.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-resource-stats-azure-sql-database?view=azuresqldb-current>

QUESTION 41

You need to configure a runbook to execute when a Microsoft Azure SQL database uses more than 80 percent of the Database Throughput Unit (DTU) resources.

What should you configure in the runbook?

- A. A trigger
- B. A Connection
- C. A Webhook
- D. A Job

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference: