

DP-300.premium.82q - DEMO

Number: DP-300
Passing Score: 800
Time Limit: 120 min



DP-300

Administering Relational Databases on Microsoft Azure



Testlet 1

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office is separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
 - Minimize downtime during the migration of the SERVER1 databases.
 - Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- Once all requirements are met, minimize costs whenever possible.

QUESTION 1

HOTSPOT

You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.

What should you include in the migration plan? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Database Migration service

Box 1: Premium 4-VCore

Scenario: Migrate the SERVER1 databases to the Azure SQL Database platform.

- Minimize downtime during the migration of the SERVER1 databases.

Premium 4-vCore is for large or business critical workloads. It supports online migrations, offline migrations, and faster migration speeds.

Incorrect Answers:

The Standard pricing tier suits most small- to medium- business workloads, but it supports offline migration only.

Box 2: A VPN gateway

You need to create a Microsoft Azure Virtual Network for the Azure Database Migration Service by using the Azure Resource Manager deployment model, which provides site-to-site connectivity to your on-premises source servers by using either ExpressRoute or VPN.

Reference: <https://azure.microsoft.com/pricing/details/database-migration/>

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online>

QUESTION 2

HOTSPOT

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

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Purchasing model:

	▼
Azure virtual machine reserved instances	
DTU	
vCore	

Deployment option:

	▼
An Azure SQL Database elastic pool	
An Azure SQL Database managed instance	
A SQL Server Always On availability group	

Correct Answer:

Answer Area

Purchasing model:	<div><div></div><div>Azure virtual machine reserved instances</div><div>DTU</div><div>vCore</div></div>
Deployment option:	<div><div></div><div>An Azure SQL Database elastic pool</div><div>An Azure SQL Database managed instance</div><div>A SQL Server Always On availability group</div></div>

Section: [none]
Explanation

Explanation/Reference:
Explanation:

Box 1: DTU

Scenario:

- The 30 new databases must scale automatically.
- Once all requirements are met, minimize costs whenever possible.

You can configure resources for the pool based either on the DTU-based purchasing model or the vCore-based purchasing model.

In short, for simplicity, the DTU model has an advantage. Plus, if you're just getting started with Azure SQL Database, the DTU model offers more options at the lower end of performance, so you can get started at a lower price point than with vCore.

Box 2: An Azure SQL database elastic pool

Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases that have varying and unpredictable usage demands. The databases in an elastic pool are on a single server and share a set number of resources at a set price. Elastic pools in Azure SQL Database enable SaaS developers to optimize the price performance for a group of databases within a prescribed budget while delivering performance elasticity for each database.

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

Question Set 2**QUESTION 1**

You have a Microsoft SQL Server 2019 instance in an on-premises datacenter. The instance contains a 4-TB database named DB1.

You plan to migrate DB1 to an Azure SQL Database managed instance.

What should you use to minimize downtime and data loss during the migration?

- A. distributed availability groups
- B. database mirroring
- C. log shipping
- D. Database Migration Assistant

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The Data Migration Assistant (DMA) helps you upgrade to a modern data platform by detecting compatibility issues that can impact database functionality in your new version of SQL Server or Azure SQL Database. DMA recommends performance and reliability improvements for your target environment and allows you to move your schema, data, and uncontained objects from your source server to your target server.

Note: SQL Managed Instance supports the following database migration options (currently these are the only supported migration methods):

- Azure Database Migration Service - migration with near-zero downtime.
- Native RESTORE DATABASE FROM URL - uses native backups from SQL Server and requires some downtime.

Reference: <https://docs.microsoft.com/en-us/sql/dma/dma-overview>

QUESTION 2

HOTSPOT

You have an on-premises Microsoft SQL Server 2016 server named Server1 that contains a database named DB1.

You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.

How should you configure the backup of DB1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Backup type:

	▼
Full and log backups only	
Full backup only	
Log backup only	

Backup option:

	▼
WITH CHECKSUM	
WITH NOINIT	
WITH UNLOAD	

Correct Answer:

Answer Area

Backup type:

	▼
Full and log backups only	
Full backup only	
Log backup only	

Backup option:

	▼
WITH CHECKSUM	
WITH NOINIT	
WITH UNLOAD	

Section: [none]**Explanation****Explanation/Reference:**

Explanation:

Box 1: Full and log backups only

Make sure to take every backup on a separate backup media (backup files). Azure Database Migration Service doesn't support backups that are appended to a single backup file. Take full backup and log backups to separate backup files.

Box 2: WITH CHECKSUM Azure Database Migration Service uses the backup and restore method to migrate your on-premises databases to SQL Managed Instance. Azure Database Migration Service only supports backups created using checksum.

Incorrect Answers:

NOINIT Indicates that the backup set is appended to the specified media set, preserving existing backup sets. If a media password is defined for the media set, the password must be supplied. NOINIT is the default.

UNLOAD

Specifies that the tape is automatically rewound and unloaded when the backup is finished. UNLOAD is the default when a session begins.

Reference: <https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online>**QUESTION 3**

DRAG DROP

You have a resource group named App1Dev that contains an Azure SQL Database server named DevServer1. DevServer1 contains an Azure SQL database named DB1. The schema and permissions for DB1 are saved in a Microsoft SQL Server Data Tools (SSDT) database project.

You need to populate a new resource group named App1Test with the DB1 database and an Azure SQL Server named TestServer1. The resources in App1Test must have the same configurations as the resources in App1Dev.

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

Select and Place:**Correct Answer:****Section: [none]****Explanation****Explanation/Reference:****QUESTION 4** You have 20 Azure SQL databases provisioned by using the vCore purchasing model.

You plan to create an Azure SQL Database elastic pool and add the 20 databases.

Which three metrics should you use to size the elastic pool to meet the demands of your workload? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. total size of all the databases
- B. geo-replication support
- C. number of concurrently peaking databases * peak CPU utilization per database
- D. maximum number of concurrent sessions for all the databases
- E. total number of databases * average CPU utilization per database

Correct Answer: ACE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

CE: Estimate the vCores needed for the pool as follows:

For vCore-based purchasing model: $\text{MAX}(\text{Total number of DBs} \times \text{average vCore utilization per DB}, \text{Number of concurrently peaking DBs} \times \text{Peak vCore utilization per DB})$ A:

Estimate the storage space needed for the pool by adding the number of bytes needed for all the databases in the pool.

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

QUESTION 5

DRAG DROP

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

A recent failure causes the master database to be lost.

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

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

Select and Place:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Step 1: Attach the SSISDB database

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

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

Step 3: Open the master key for the SSISDB database

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

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

Alter Master Key Add encryption by Service Master Key

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

Reference: <https://docs.microsoft.com/en-us/sql/integration-services/backup-restore-and-move-the-ssis-catalog>

QUESTION 6

You have an Azure SQL database that contains a table named factSales. FactSales contains the columns shown in the following table.

Name	Data type
SalesID	Int
Product	Int
Total Number	Numeric(8,4)
Tax Number	Numeric(8,4)
SalesRep	Varchar(30)

FactSales has 6 billion rows and is loaded nightly by using a batch process.

Which type of compression provides the greatest space reduction for the database?

- A. page compression
- B. row compression
- C. columnstore compression
- D. columnstore archival compression

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Columnstore tables and indexes are always stored with columnstore compression. You can further reduce the size of columnstore data by configuring an additional compression called archival compression.

Note: Columnstore — The columnstore index is also logically organized as a table with rows and columns, but the data is physically stored in a column-wise data format.

Incorrect Answers:

B: Rowstore — The rowstore index is the traditional style that has been around since the initial release of SQL Server.

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression>

QUESTION 7

You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

- Clustered columnstore indexes
 - Automatic tuning
 - Change tracking ▪
- PolyBase

You plan to migrate DB1 to an Azure SQL database.

What feature should be removed or replaced before DB1 can be migrated?

- A. Clustered columnstore indexes
- B. PolyBase
- C. Change tracking
- D. Automatic tuning

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

This table lists the key features for PolyBase and the products in which they're available.

Feature	SQL Server (Beginning with 2016)	Azure SQL Database	Azure Synapse Analytics	Parallel Data Warehouse
Query Hadoop data with Transact-SQL	Yes	No	No	Yes
Import data from Hadoop	Yes	No	No	Yes
Export data to Hadoop	Yes	No	No	Yes
Query, import from, export to Azure HDInsight	No	No	No	No
Push down query computations to Hadoop	Yes	No	No	Yes
Import data from Azure Blob storage	Yes	Yes*	Yes	Yes
Export data to Azure Blob storage	Yes	No	Yes	Yes
Import data from Azure Data Lake Store	No	No	Yes	No
Export data to Azure Data Lake Store	No	No	Yes	No
Run PolyBase queries from Microsoft BI tools	Yes	No	Yes	Yes

VCEUp

Incorrect Answers:

C: Change tracking is a lightweight solution that provides an efficient change tracking mechanism for applications. It applies to both Azure SQL Database and SQL Server.

D: Azure SQL Database and Azure SQL Managed Instance automatic tuning provides peak performance and stable workloads through continuous performance tuning based on AI and machine learning.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary>

Testlet 1

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises server named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contains database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1.

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints. **Business**

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
 - Minimize downtime during the migration of the SERVER1 databases.
 - Use the Azure Hybrid Use Benefits when migrating workloads to Azure. ▪
- Once all requirements are met, minimize costs whenever possible.

QUESTION 1

DRAG DROP

You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance 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:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

1. Create and populate Azure AD.
2. Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
3. Create an Azure Active Directory administrator. (Step 1)
4. Configure your client computers.
5. Create contained database users in your database mapped to Azure AD identities. (Step 2)
6. Connect to your database by using Azure AD identities. (Step 3)

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

QUESTION 2

DRAG DROP

You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance 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**

- Modify the Azure SQL server administrator account.
- Create database users.
- Create a user in the master database.
- Create an Azure AD administrator for the logical server.
- Connect to the databases by using an Azure AD account.



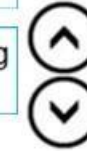
Correct Answer:

Actions**Answer Area**

- Modify the Azure SQL server administrator account.
-
- Create a user in the master database.
-
-



- Create an Azure AD administrator for the logical server.
- Create database users.
- Connect to the databases by using an Azure AD account.



Section: [none]
Explanation

Explanation/Reference:
 Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

1. Create and populate Azure AD.
2. Optional: Associate or change the active directory that is currently associated with your Azure Subscription.
3. Create an Azure Active Directory administrator. (Step 1)
4. Configure your client computers.
5. Create contained database users in your database mapped to Azure AD identities. (Step 2)
6. Connect to your database by using Azure AD identities. (Step 3)

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

Testlet 2

Case study

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

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

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

To start the case study

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

Overview

Existing Environment

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

Active Directory

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

Database Environment

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

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

VCEUp

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

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

Current Business Model

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

Requirements

Planned Changes

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

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

- Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.
 - Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months. ▪
- Start onboarding customers to the new PaaS solution within six months.

Business Goals

Contoso identifies the following business requirements:

- Use built-in Azure features whenever possible.
- Minimize development effort whenever possible.
- Minimize the compute costs of the PaaS solutions.
- Provide all the customers with their own copy of the database by using the PaaS solution.
- Provide the customers with different table and row access based on the customer's service agreement.
- In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.
- Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso. **Technical**

Requirements

Contoso identifies the following technical requirements:

- Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure AD of Contoso to minimize guest users.
- All customers must have their own resource group, Azure SQL server, and Azure SQL database. The deployment of resources for each customer must be done in a consistent fashion.
- Users must be able to review the queries issued against the PaaS databases and identify any new objects created.
- Downtime during the PostgreSQL database migration must be minimized.

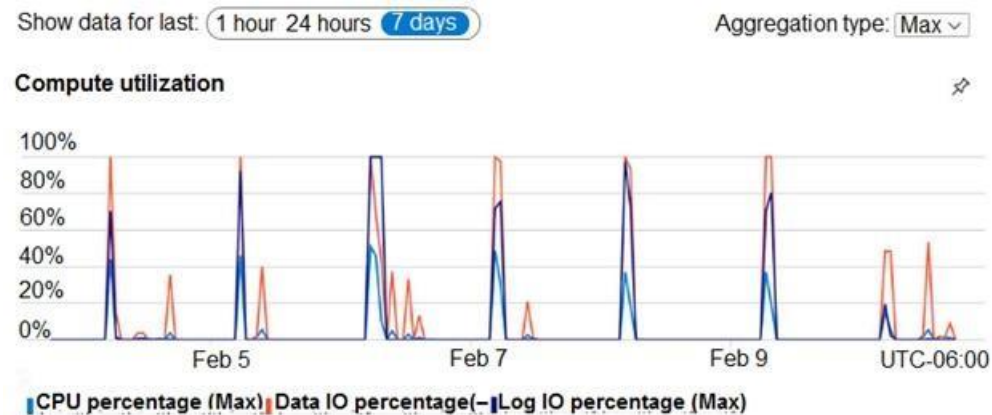
Monitoring Requirements

Contoso identifies the following monitoring requirements:

- Notify administrators when a PaaS database has a higher than average CPU usage.
- Use a single dashboard to review security and audit data for all the PaaS databases.
- Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.
- Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues automatically whenever possible.

PaaS Prototype

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



Role Assignments

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

QUESTION 1

HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Hot Area:

Statements

Yes

No

☐☐☐☐

IT Exam Dumps – Learn Anything | VCEup.com

Answer Area

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

Section: [none]
Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

DBAGroup1 is member of the Contributor role.

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

Box 2: No

Box 3: Yes

DBAGroup2 is member of the SQL DB Contributor role.

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

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

Question Set 3**QUESTION 1**

You have a new Azure SQL database. The database contains a column that stores confidential information.

You need to track each time values from the column are returned in a query. The tracking information must be stored for 365 days from the date the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Turn on auditing and write audit logs to an Azure Storage account.
- B. Add extended properties to the column.
- C. Turn on Advanced Data Security for the Azure SQL server.
- D. Apply sensitivity labels named Highly Confidential to the column.
- E. Turn on Azure Advanced Threat Protection (ATP).

Correct Answer: ACD

Section: [none]
Explanation

Explanation/Reference:

Explanation:

C: Advanced Data Security (ADS) is a unified package for advanced SQL security capabilities. ADS is available for Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics. It includes functionality for discovering and classifying sensitive data

D: You can apply sensitivity-classification labels persistently to columns by using new metadata attributes that have been added to the SQL Server database engine. This metadata can then be used for advanced, sensitivity-based auditing and protection scenarios.

A: An important aspect of the information-protection paradigm is the ability to monitor access to sensitive data. Azure SQL Auditing has been enhanced to include a new field in the audit log called data_sensitivity_information. This field logs the sensitivity classifications (labels) of the data that was returned by a query. Here's an example:

d	client_ip	application_name	duration_milliseconds	response_rows	affected_rows	connection_id	data_sensitivity_information
	7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	7.125	Microsoft SQL Server Management Studio - Query	41	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/data-discovery-and-classification-overview>

QUESTION 2

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

- Ensure that all traffic to the public endpoint of SqlSrv1 is blocked.
- Minimize the possibility of VM1 exfiltrating data stored in SqlDb1.

What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link
- D. an ExpressRoute gateway

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

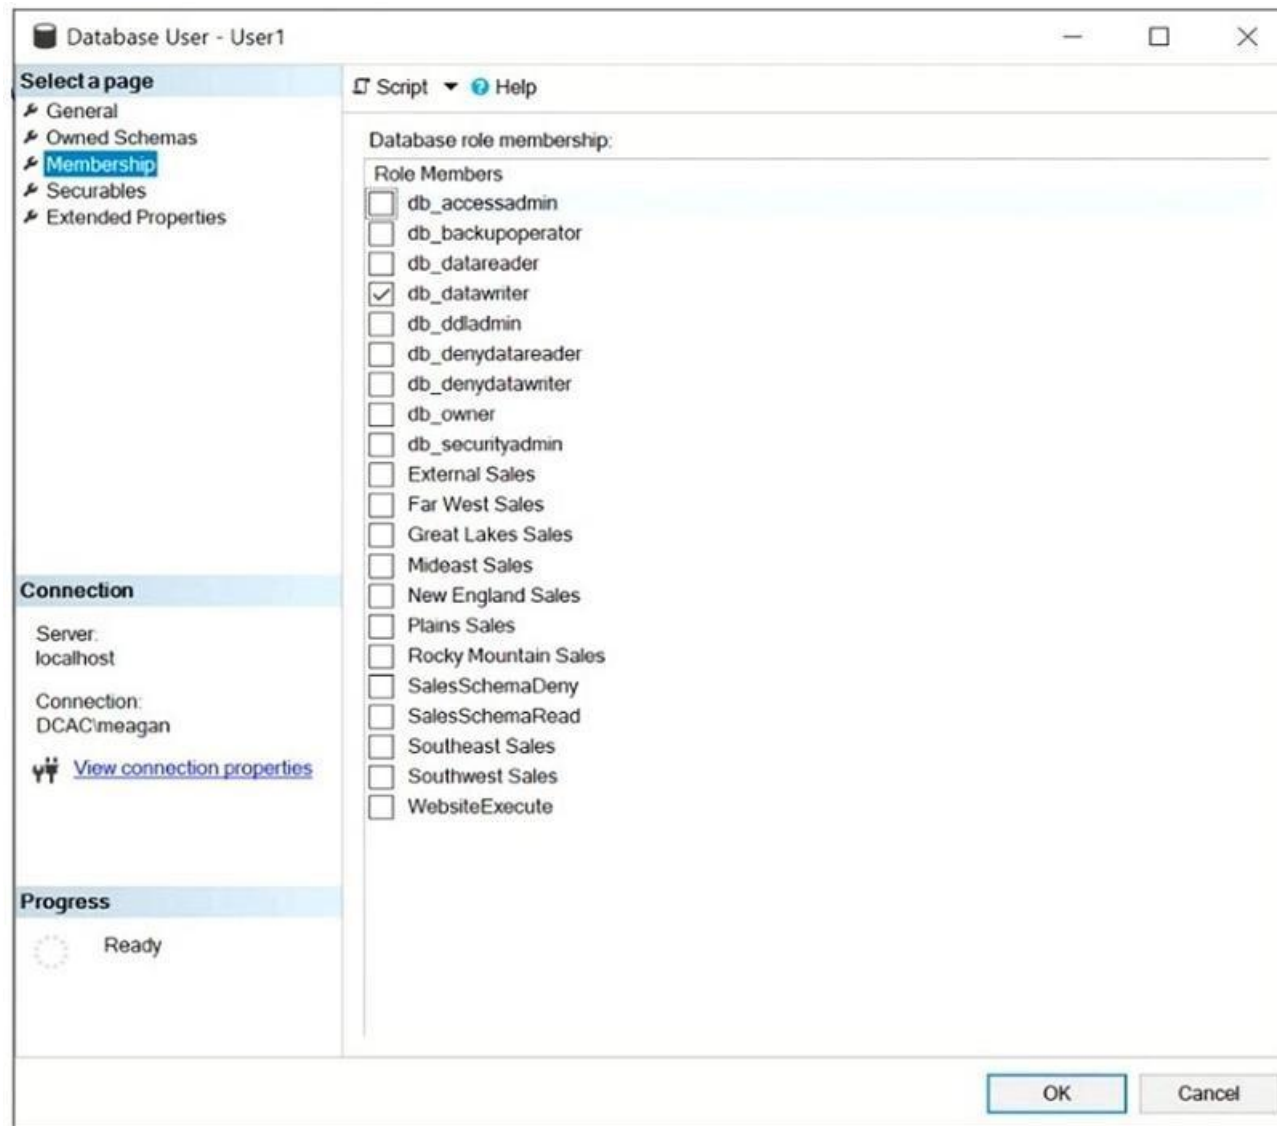
Reference: <https://docs.microsoft.com/en-us/azure/private-link/private-link-overview>

QUESTION 3

HOTSPOT

You have a Microsoft SQL Server database named DB1 that contains a table named Table1.

The database role membership for a user named User1 is shown in the following exhibit.



VCEUp

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

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: delete a row from Table1

Members of the db_datawriter fixed database role can add, delete, or change data in all user tables.

Box 2: db_datareader

Members of the db_datareader fixed database role can read all data from all user tables.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

QUESTION 4

DRAG DROP

You have a new Azure SQL database named DB1 on an Azure SQL server named AzSQL1.

The only user who was created is the server administrator.

You need to create a contained database user in DB1 who will use Azure Active Directory (Azure AD) for authentication.

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:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Step 1: Set up the Active Directory Admin for AzSQL1.

Step 2: Connect to DB1 by using the server administrator.

Sign into your managed instance with an Azure AD login granted with the sysadmin role.

Step 3: Create a user by using the FROM EXTERNAL PROVIDER clause.

FROM EXTERNAL PROVIDER is available for creating server-level Azure AD logins in SQL Database managed instance. Azure AD logins allow database-level Azure AD principals to be mapped to server-level Azure AD logins. To create an Azure AD user from an Azure AD login use the following syntax:

CREATE USER [AAD_principal] FROM LOGIN [Azure AD login]

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

VCEUp

QUESTION 5

HOTSPOT

You have an Azure SQL database that contains a table named Customer. Customer has the columns shown in the following table.

Customer_ID	Customer_Name	Customer_Phone
11001	Contoso, Ltd.	555-555-0173
11002	Litware, Inc.	555-505-3124
11003	ADatum Corporation	555-689-4312

You plan to implement a dynamic data mask for the Customer_Phone column. The mask must meet the following requirements:

- The first six numerals of each customer's phone number must be masked.
- The last four digits of each customer's phone number must be visible. ▪

Hyphens must be preserved and displayed.

How should you configure the dynamic data mask? To answer, select the appropriate options in the answer area.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: 0

Custom String : Masking method that exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix

Box 2: xxx-xxx

Box 3: 5

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/dynamic-data-masking>

QUESTION 6

DRAG DROP

You have an Azure SQL database that contains a table named Employees. Employees contains a column named Salary.

You need to encrypt the Salary column. The solution must prevent database administrators from reading the data in the Salary column and must provide the most secure encryption.

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:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Step 1: Create a column master key

Create a column master key metadata entry before you create a column encryption key metadata entry in the database and before any column in the database can be encrypted using Always Encrypted.

Step 2: Create a column encryption key.

Step 3: Encrypt the Salary column by using the randomized encryption type.

Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.

Note: A column encryption key metadata object contains one or two encrypted values of a column encryption key that is used to encrypt data in a column. Each value is encrypted using a column master key.

Incorrect Answers:

Deterministic encryption.

Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns. However, it may also allow unauthorized users to guess information about encrypted values by examining patterns in the encrypted column, especially if there's a small set of possible encrypted values, such as True/False, or North/South/East/West region.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

QUESTION 7

HOTSPOT

You have an Azure SQL database named DB1 that contains two tables named Table1 and Table2. Both tables contain a column named a Column1. Column1 is used for joins by an application named App1.

You need to protect the contents of Column1 at rest, in transit, and in use.

How should you protect the contents of Column1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Encryption key:

	▼
Column encryption key	
Database encryption key	
Service master key	

Encryption type:

	▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	

Correct Answer:

Answer Area

Encryption key:

	▼
Column encryption key	
Database encryption key	
Service master key	

Encryption type:

	▼
Deterministic	
Randomized	
Transparent Data Encryption (TDE)	

Section: [none]

Explanation**Explanation/Reference:**

Explanation:

Box 1: Column encryption Key

Always Encrypted uses two types of keys: column encryption keys and column master keys. A column encryption key is used to encrypt data in an encrypted column. A column master key is a key-protecting key that encrypts one or more column encryption keys.

Incorrect Answers:

TDE encrypts the storage of an entire database by using a symmetric key called the Database Encryption Key (DEK).

Box 2: Deterministic

Always Encrypted is a feature designed to protect sensitive data, such as credit card numbers or national identification numbers (for example, U.S. social security numbers), stored in Azure SQL Database or SQL Server databases. Always Encrypted allows clients to encrypt sensitive data inside client applications and never reveal the encryption keys to the Database Engine (SQL Database or SQL Server).

Always Encrypted supports two types of encryption: randomized encryption and deterministic encryption.

Deterministic encryption always generates the same encrypted value for any given plain text value. Using deterministic encryption allows point lookups, equality joins, grouping and indexing on encrypted columns.

Incorrect Answers:

- Randomized encryption uses a method that encrypts data in a less predictable manner. Randomized encryption is more secure, but prevents searching, grouping, indexing, and joining on encrypted columns.
- Transparent data encryption (TDE) helps protect Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics against the threat of malicious offline activity by encrypting data at rest. It performs real-time encryption and decryption of the database, associated backups, and transaction log files at rest without requiring changes to the application.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

QUESTION 8

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

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

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

NOTE: Each correct selection is worth one point.

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

Correct Answer: BC

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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

Database-level firewall rules only apply to individual databases.

Incorrect Answers:

B: Server-level IP firewall rules apply to all databases within the same server.

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

QUESTION 9

DRAG DROP

You have an Azure SQL Database instance named DatabaseA on a server named Server1.

You plan to add a new user named App1 to DatabaseA and grant App1 db_datacenter permissions. App1 will use SQL Server Authentication.

You need to create App1. The solution must ensure that App1 can be given access to other databases by using the same credentials.

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**

On the master database, run CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;

On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';

On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];

On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';

On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];



Correct Answer:

Actions**Answer Area**

On the master database, run CREATE LOGIN [APP1] FROM EXTERNAL PROVIDER;

On DatabaseA, run CREATE USER [APP1] WITH PASSWORD = 'P@ssW0rd!';



On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'P@aaW0rd!';

On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1];

On DatabaseA, run ALTER ROLE db_datareader ADD MEMBER [App1];



Section: [none]

Explanation

Explanation/Reference:

Explanation:

Step 1: On the master database, run CREATE LOGIN [App1] WITH PASSWORD = 'p@aaW0rd!'

Logins are server wide login and password pairs, where the login has the same password across all databases. Here is some sample Transact-SQL that creates a login:

```
CREATE LOGIN readonlylogin WITH password='1231!#ASDF!a';
```

You must be connected to the master database on SQL Azure with the administrative login (which you get from the SQL Azure portal) to execute the CREATE LOGIN command.

Step 2: On DatabaseA, run CREATE USER [App1] FROM LOGIN [App1]

Users are created per database and are associated with logins. You must be connected to the database in where you want to create the user. In most cases, this is not the master database. Here is some sample Transact-SQL that creates a user:

```
CREATE USER readonlyuser FROM LOGIN readonlylogin;
```

Step 3: On DatabaseA run ALTER ROLE db_datareader ADD Member [App1]

Just creating the user does not give them permissions to the database. You have to grant them access. In the Transact-SQL example below the readonlyuser is given read only permissions to the database via the db_datareader role. EXEC

```
sp_addrolemember 'db_datareader', 'readonlyuser';
```

Reference:

<https://azure.microsoft.com/en-us/blog/adding-users-to-your-sql-azure-database/>

QUESTION 10

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1. Restrict network connectivity to SqlSrv1.

What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private link
- D. an ExpressRoute gateway

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.

Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary.

Reference:

<https://docs.microsoft.com/en-us/azure/private-link/private-link-overview>

Testlet 1

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.

- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints. **Business**

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
- Minimize downtime during the migration of the SERVER1 databases.
- Use the Azure Hybrid Use Benefits when migrating workloads to Azure.

Once all requirements are met, minimize costs whenever possible.

QUESTION 1

HOTSPOT

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: The server, the elastic pool, and the database

Scenario:

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Box 2: Azure Event hubs

Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform. Event hubs are able to handle custom metrics.

Incorrect Answers:

Azure Log Analytics

Azure metric and log data are sent to Azure Monitor Logs, previously known as Azure Log Analytics, directly by Azure. Azure SQL Analytics is a cloud only monitoring solution supporting streaming of diagnostics telemetry for all of your Azure SQL databases.

However, because Azure SQL Analytics does not use agents to connect to Azure Monitor, it does not support monitoring of SQL Server hosted on-premises or in virtual machines.

Testlet 2

Case study

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

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

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

To start the case study

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

Overview

Existing Environment

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

Active Directory

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

Database Environment

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

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

VCEUp

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

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

Current Business Model

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

Requirements

Planned Changes

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

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

- Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.
 - Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months. ▪
- Start onboarding customers to the new PaaS solution within six months.

Business Goals

Contoso identifies the following business requirements:

- Use built-in Azure features whenever possible.
- Minimize development effort whenever possible.
- Minimize the compute costs of the PaaS solutions.
- Provide all the customers with their own copy of the database by using the PaaS solution.
- Provide the customers with different table and row access based on the customer's service agreement.
- In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.
- Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso. **Technical**

Requirements

Contoso identifies the following technical requirements:

- Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure AD of Contoso to minimize guest users.
- All customers must have their own resource group, Azure SQL server, and Azure SQL database. The deployment of resources for each customer must be done in a consistent fashion.
- Users must be able to review the queries issued against the PaaS databases and identify any new objects created.
- Downtime during the PostgreSQL database migration must be minimized.

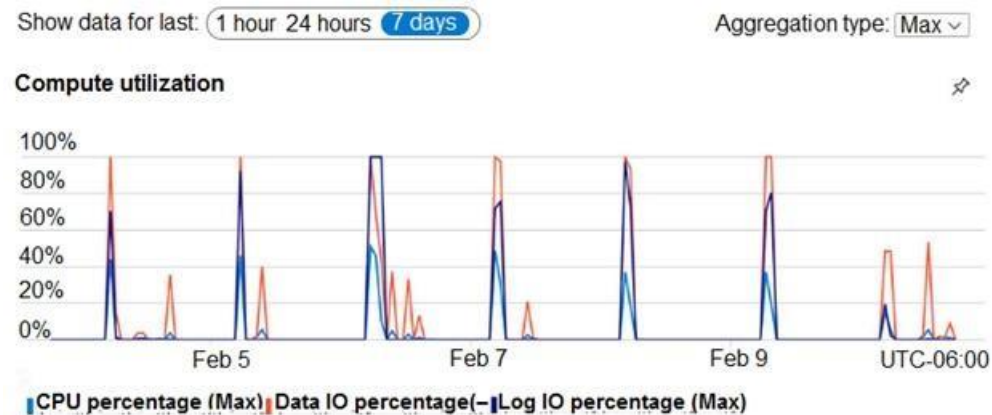
Monitoring Requirements

Contoso identifies the following monitoring requirements:

- Notify administrators when a PaaS database has a higher than average CPU usage.
- Use a single dashboard to review security and audit data for all the PaaS databases.
- Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.
- Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues automatically whenever possible.

PaaS Prototype

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



Role Assignments

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

+

 Add

≡

 Edit columns

↺

 Refresh

🗑

 Remove

♥

 Got feedback?

[Check access](#)
[Role assignments](#)
[Deny assignments](#)
[Classic administrators](#)
[Roles](#)

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

Number of role assignments for this subscription ⓘ

15

2000

Name ⓘ

Search by name or email

Type ⓘ

Groups

Role ⓘ

2 selected

Scope ⓘ

All scopes



Group by ⓘ

Role

ⓘ

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

2 items (2 Groups)

<input type="checkbox"/>	Name	Type	Role	Scope
Contributor				
<input type="checkbox"/>	 DBAGroup1	Group	Contributor ⓘ	This resource
SQL DB Contributor				
<input type="checkbox"/>	 DBAGroup2	Group	SQL DB Contributor ⓘ	This resource

QUESTION 1 Based on the PaaS prototype, which Azure SQL Database compute tier should you use?

- A. Business Critical 4-vCore
- B. Hyperscale
- C. General Purpose v-vCore
- D. Serverless

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

There are CPU and Data I/O spikes for the PaaS prototype. Business Critical 4-vCore is needed.

Incorrect Answers: B: Hyperscale is for large databases

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

QUESTION 2 Which audit log destination should you use to meet the monitoring requirements?

- A. Azure Storage
- B. Azure Event Hubs
- C. Azure Log Analytics

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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

With dashboards can bring together operational data that is most important to IT across all your Azure resources, including telemetry from Azure Log Analytics.

Note: Auditing for Azure SQL Database and Azure Synapse Analytics tracks database events and writes them to an audit log in your Azure storage account, Log Analytics workspace, or Event Hubs.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/visualize/tutorial-logs-dashboards>

Question Set 3**QUESTION 1**

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

You need to record whenever users query the CustomerPII table.

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

NOTE: Each correct selection is worth one point.

- A. server audit specification
- B. SQL Server audit
- C. database audit specification
- D. a server principal

Correct Answer: AC

Section: [none]

Explanation

Explanation/Reference:

Explanation:

An auditing policy can be defined for a specific database or as a default server policy in Azure (which hosts SQL Database or Azure Synapse):

- A server policy applies to all existing and newly created databases on the server.
- If server auditing is enabled, it always applies to the database. The database will be audited, regardless of the database auditing settings.
- Enabling auditing on the database, in addition to enabling it on the server, does not override or change any of the settings of the server auditing. Both audits will exist side by side.

Note:

The Server Audit Specification object belongs to an audit.

A Database Audit Specification defines which Audit Action Groups will be audited for the specific database in which the specification is created.

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

QUESTION 2 You have an Azure virtual machine based on a custom image named VM1.

VM1 hosts an instance of Microsoft SQL Server 2019 Standard.

You need to automate the maintenance of VM1 to meet the following requirements:

- Automate the patching of SQL Server and Windows Server.
- Automate full database backups and transaction log backups of the databases on VM1.
- Minimize administrative effort.

What should you do first?

- A. Enable a system-assigned managed identity for VM1
- B. Register VM1 to the Microsoft.Sql resource provider
- C. Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1D. Register VM1 to the Microsoft.SqlVirtualMachine resource provider

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Automated Patching depends on the SQL Server infrastructure as a service (IaaS) Agent Extension. The SQL Server IaaS Agent Extension (SqlIaaSExtension) runs on Azure virtual machines to automate administration tasks. The SQL Server IaaS extension is installed when you register your SQL Server VM with the SQL Server VM resource provider.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/sql-server-iaas-agent-extension-automate-management>

QUESTION 3 You receive numerous alerts from Azure Monitor for an Azure SQL database.

You need to reduce the number of alerts. You must only receive alerts if there is a significant change in usage patterns for an extended period.

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

NOTE: Each correct selection is worth one point.

- A. Set Threshold Sensitivity to **High**
- B. Set the Alert logic threshold to **Dynamic**
- C. Set the Alert logic threshold to **Static**
- D. Set Threshold Sensitivity to **Low**
- E. Set Force Plan to **On**

Correct Answer: BD

Section: [none]

Explanation

Explanation/Reference:

Explanation:

B: Dynamic Thresholds continuously learns the data of the metric series and tries to model it using a set of algorithms and methods. It detects patterns in the data such as seasonality (Hourly / Daily / Weekly), and is able to handle noisy metrics (such as machine CPU or memory) as well as metrics with low dispersion (such as availability and error rate).

D: Alert threshold sensitivity is a high-level concept that controls the amount of deviation from metric behavior required to trigger an alert.

Low – The thresholds will be loose with more distance from metric series pattern. An alert rule will only trigger on large deviations, resulting in fewer alerts.

Incorrect Answers:

A: High – The thresholds will be tight and close to the metric series pattern. An alert rule will be triggered on the smallest deviation, resulting in more alerts.

Reference: <https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-dynamic-thresholds>

QUESTION 4

You have an Azure SQL database named sqldb1.

You need to minimize the amount of space by the data and log files of sqldb1.

What should you run?

- A. DBCC SHRINKDATABASE
- B. sp_clean_db_free_space
- C. sp_clean_db_file_free_space
- D. DBCC SHRINKFILE

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

DBCC SHRINKDATABASE shrinks the size of the data and log files in the specified database.

Incorrect Answers:

D: To shrink one data or log file at a time for a specific database, execute the DBCC SHRINKFILE command.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-shrinkdatabase-transact-sql>

QUESTION 5 You have an Azure SQL Database server named sqlsrv1 that hosts 10 Azure SQL databases.

The databases perform slower than expected.

You need to identify whether the performance issue relates to the use of tempdb by Azure SQL databases on sqlsrv1.

What should you do?

- A. Run Query Store-based queries
- B. Review information provided by SQL Server Profiler-based traces
- C. Review information provided by Query Performance Insight
- D. Run dynamic management view-based queries

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting.

You can use the Intelligent Insights performance diagnostics log of Azure SQL Database to troubleshoot performance issues.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb-contention> <https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-use-diagnostics-log>

QUESTION 6

DRAG DROP

You are building an Azure virtual machine.

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

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

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

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

Select and Place:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Follow these same steps to create striped virtual disk:

- Create Log Storage Pool.
- Create Virtual Disk
- Create Volume

Box 1: a storage pool

Box 2: a virtual disk that uses stripe layout

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

Box 3: a volume

Reference: <https://hanu.com/hanu-how-to-striping-of-disks-for-azure-sql-server/>

QUESTION 7

You have an Azure SQL database named sqldb1.

You need to minimize the possibility of Query Store transitioning to a read-only state.

What should you do?

- A. Double the value of Data Flush interval
- B. Decrease by half the value of Data Flush Interval
- C. Double the value of Statistics Collection Interval
- D. Decrease by half the value of Statistics Collection interval

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The Max Size (MB) limit isn't strictly enforced. Storage size is checked only when Query Store writes data to disk. This interval is set by the Data Flush Interval (Minutes) option. If Query Store has breached the maximum size limit between storage size checks, it transitions to read-only mode.

Incorrect Answers:

C: Statistics Collection Interval: Defines the level of granularity for the collected runtime statistic, expressed in minutes. The default is 60 minutes. Consider using a lower value if you require finer granularity or less time to detect and mitigate issues. Keep in mind that the value directly affects the size of Query Store data.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/best-practice-with-the-query-store>

QUESTION 8

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.

You scale up the virtual machine to 16 vCPUs and 64 GB of memory.

You need to provide the lowest latency for tempdb.

What is the total number of data files that tempdb should contain?

- A. 2
- B. 4
- C. 8
- D. 64

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

QUESTION 9

HOTSPOT

You have an Azure SQL database named db1.

You need to retrieve the resource usage of db1 from the last week.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: sys.resource_stats sys.resource_stats returns CPU usage and storage data for an Azure SQL Database. It has database_name and start_time columns.

Box 2: DateAdd

The following example returns all databases that are averaging at least 80% of compute utilization over the last one week.

```
DECLARE @s datetime;
DECLARE @e datetime;
SET @s= DateAdd(d,-7,GetUTCDate());
SET @e= GETUTCDATE();
SELECT database_name, AVG(avg_cpu_percent) AS Average_Compute_Utilization
FROM sys.resource_stats
WHERE start_time BETWEEN @s AND @e
GROUP BY database_name
HAVING AVG(avg_cpu_percent) >= 80
```

Incorrect Answers:

sys.dm_exec_requests: sys.dm_exec_requests returns information about each request that is executing in SQL Server. It does not have a column named database_name.

sys.dm_db_resource_stats: sys.dm_db_resource_stats does not have any start_time column.

Note: sys.dm_db_resource_stats 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. Historical data is maintained for approximately one hour.

Sys.dm_user_db_resource_governance returns actual configuration and capacity settings used by resource governance mechanisms in the current database or elastic pool. It does not have any start_time column.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-resource-stats-azure-sql-database>

QUESTION 10

You have 50 Azure SQL databases.

You need to notify the database owner when the database settings, such as the database size and pricing tier, are modified in Azure.

What should you do?

- A. Create a diagnostic setting for the activity log that has the Security log enabled.
- B. For the database, create a diagnostic setting that has the InstanceAndAppAdvanced metric enabled.
- C. Create an alert rule that uses a Metric signal type.
- D. Create an alert rule that uses an Activity Log signal type.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Activity log events - An alert can trigger on every event, or, only when a certain number of events occur.

Incorrect Answers:

C: Metric values - The alert triggers when the value of a specified metric crosses a threshold you assign in either direction. That is, it triggers both when the condition is first met and then afterwards when that condition is no longer being met.

Reference:

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

QUESTION 11

You have several Azure SQL databases on the same Azure SQL Database server in a resource group named ResourceGroup1.

You must be alerted when CPU usage exceeds 80 percent for any database. The solution must apply to any additional databases that are created on the Azure SQL server.

Which resource type should you use to create the alert?

- A. Resource Groups
- B. SQL Servers
- C. SQL Databases
- D. SQL Virtual Machines

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

There are resource types related to application code, compute infrastructure, networking, storage + databases.

You can deploy up to 800 instances of a resource type in each resource group.

Some resources can exist outside of a resource group. These resources are deployed to the subscription, management group, or tenant. Only specific resource types are supported at these scopes.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/resource-providers-and-types>

QUESTION 12

You have SQL Server 2019 on an Azure virtual machine that runs Windows Server 2019. The virtual machine has 4 vCPUs and 28 GB of memory.

You scale up the virtual machine to 8 vCPUs and 64 GB of memory.

You need to provide the lowest latency for tempdb.

What is the total number of data files that tempdb should contain?

- A. 2
- B. 4
- C. 8
- D. 64

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

Testlet 1

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises server named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contains database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1.

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints. **Business**

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
 - Minimize downtime during the migration of the SERVER1 databases.
 - Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- Once all requirements are met, minimize costs whenever possible.

QUESTION 1 You need to identify the cause of the performance issues on SalesSQLDb1.

Which two dynamic management views should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. `sys.dm_pdw_nodes_tran_locks`
- B. `sys.dm_exec_compute_node_errors`
- C. `sys.dm_exec_requests`
- D. `sys.dm_cdc_errors`
- E. `sys.dm_pdw_nodes_os_wait_stats`
- F. `sys.dm_tran_locks`

Correct Answer: AE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

A: Use sys.dm_pdw_nodes_tran_locks instead of sys.dm_tran_locks from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information.

```
SELECT
    t1.resource_type,
t1.resource_database_id,
t1.resource_associated_entity_id,
t1.request_mode,
t1.request_session_id,
t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
    ON t1.lock_owner_address = t2.resource_address;
```

Note: Depending on the system you're working with you can access these wait statistics from one of three locations:

sys.dm_os_wait_stats: for SQL Server sys.dm_db_wait_stats: for
Azure SQL Database sys.dm_pdw_nodes_os_wait_stats: for Azure
SQL Data Warehouse

Incorrect Answers:

F: sys.dm_tran_locks returns information about currently active lock manager resources in SQL Server 2019 (15.x). Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be granted.

Instead use sys.dm_pdw_nodes_tran_locks.

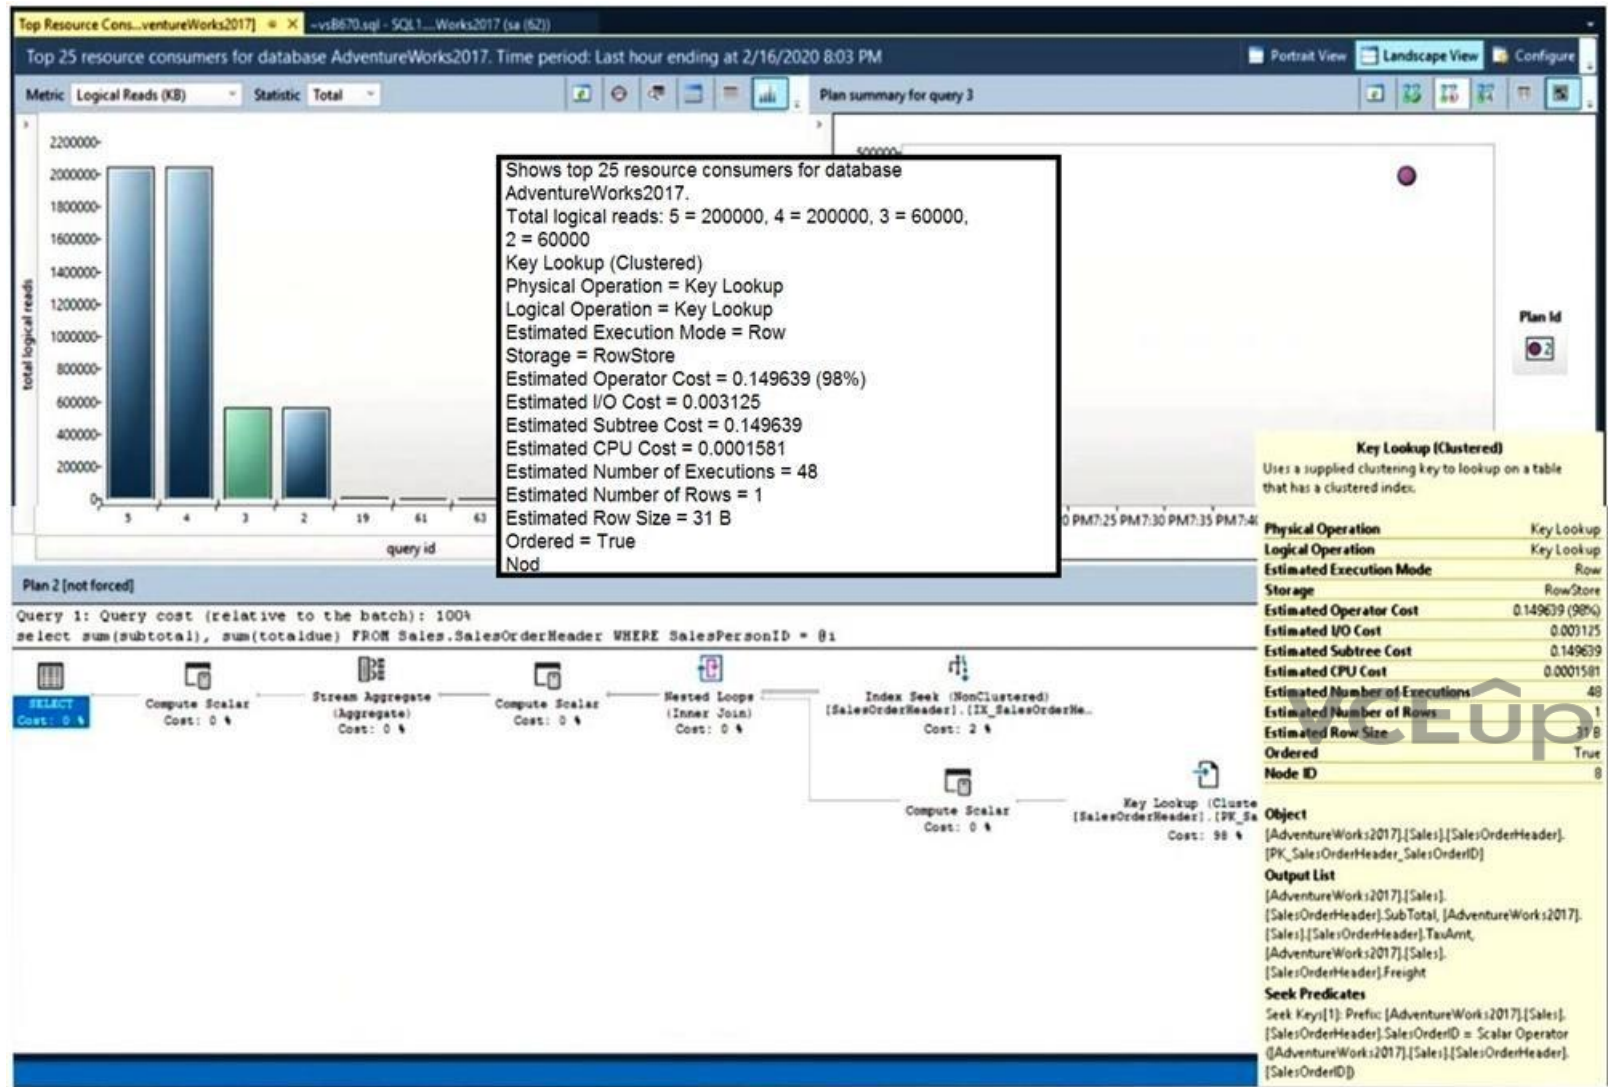
Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-locks-transact-sql>

Question Set 2

QUESTION 1
HOTSPOT

You have SQL Server on an Azure virtual machine.

You review the query plan shown in the following exhibit.



For each of the following statements, select yes if the statement is true. Otherwise, select no.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<input type="radio"/>	<input type="radio"/>
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	<input type="radio"/>	<input type="radio"/>
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	<input checked="" type="radio"/>	<input type="radio"/>
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	<input type="radio"/>	<input checked="" type="radio"/>
You will reduce the I/O usage and the query execution time if you include the SubTotal, TaxAmt, and Freight columns in the PK_SalesOrderHeader_SalesOrderID index.	<input checked="" type="radio"/>	<input type="radio"/>

Section: [none]

Explanation

Explanation/Reference:

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store>

QUESTION 2

A data engineer creates a table to store employee information for a new application. All employee names are in the US English alphabet. All addresses are locations in the United States. The data engineer uses the following statement to create the table.

```
CREATE TABLE dbo.Employee
```

```
(
    EmployeeID      INT IDENTITY(1,1) PRIMARY KEY CLUSTERED NOT NULL,
    FirstName       VARCHAR(100) NOT NULL,
    LastName        VARCHAR(100) NOT NULL,
    Title           VARCHAR(100) NULL,
    LastHireDate     DATETIME NULL,
    StreetAddress1   VARCHAR(500) NOT NULL,
    StreetAddress2   VARCHAR(500) NOT NULL,
    StreetAddress3   VARCHAR(500) NOT NULL,
    City            VARCHAR(200) NOT NULL,
    StateName       VARCHAR(20) NOT NULL,
    Salary          VARCHAR(20) NULL,
    PhoneNumber     VARCHAR(20) NOT NULL
)
```

You need to recommend changes to the data types to reduce storage and improve performance.

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

NOTE: Each correct selection is worth one point.

- A. Change `Salary` to the `money` data type.
- B. Change `PhoneNumber` to the `float` data type.
- C. Change `LastHireDate` to the `datetime2(7)` data type.
- D. Change `PhoneNumber` to the `bigint` data type.
- E. Change `LastHireDate` to the `date` data type.

Correct Answer: AE

Section: [none]

Explanation

Explanation/Reference:

QUESTION 3

You have an Azure SQL database.

You identify a long running query.

You need to identify which operation in the query is causing the performance issue.

What should you use to display the query execution plan in Microsoft SQL Server Management Studio (SSMS)?

- A. Live Query Statistics
- B. an estimated execution plan
- C. an actual execution plan
- D. Client Statistics

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

To include an execution plan for a query during execution

1. On the SQL Server Management Studio toolbar, click Database Engine Query. You can also open an existing query and display the estimated execution plan by clicking the Open File toolbar button and locating the existing query.
2. Enter the query for which you would like to display the actual execution plan.
3. On the Query menu, click Include Actual Execution Plan or click the Include Actual Execution Plan toolbar button.



Note: Actual execution plans are generated after the Transact-SQL queries or batches execute. Because of this, an actual execution plan contains runtime information, such as actual resource usage metrics and runtime warnings (if any). The execution plan that is generated displays the actual query execution plan that the SQL Server Database Engine used to execute the queries.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/display-an-actual-execution-plan>

QUESTION 4

You have a version-8.0 Azure Database for MySQL database.

You need to identify which database queries consume the most resources.

Which tool should you use?

- A. Query Store
- B. Metrics
- C. Query Performance Insight
- D. Alerts

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The Query Store feature in Azure Database for MySQL provides a way to track query performance over time. Query Store simplifies performance troubleshooting by helping you quickly find the longest running and most resource-intensive queries. Query Store automatically captures a history of queries and runtime statistics, and it retains them for your review. It separates data by time windows so that you can see database usage patterns. Data for all users, databases, and queries is stored in the mysql schema database in the Azure Database for MySQL instance.

Reference: <https://docs.microsoft.com/en-us/azure/mysql/concepts-query-store>

QUESTION 5 You have SQL Server on an Azure virtual machine that contains a database named DB1.

You have an application that queries DB1 to generate a sales report.

You need to see the parameter values from the last time the query was executed.

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

NOTE: Each correct selection is worth one point.

- A. Enable `Last_Query_Plan_Stats` in the master database
- B. Enable `Lightweight_Query_Profiling` in DB1
- C. Enable `Last_Query_Plan_Stats` in DB1
- D. Enable `Lightweight_Query_Profiling` in the master database
- E. Enable `PARAMETER_SNIFFING` in DB1

Correct Answer: AC

Section: [none]

Explanation

Explanation/Reference:

Explanation:

`Last_Query_Plan_Stats` allows you to enable or disable collection of the last query plan statistics (equivalent to an actual execution plan) in `sys.dm_exec_query_plan_stats`.

Lightweight profiling can be disabled at the database level using the `LIGHTWEIGHT_QUERY_PROFILING` database scoped configuration: `ALTER DATABASE SCOPED CONFIGURATION SET LIGHTWEIGHT_QUERY_PROFILING = OFF;`

Incorrect Answers:

E: Parameter sensitivity, also known as "parameter sniffing", refers to a process whereby SQL Server "sniffs" the current parameter values during compilation or recompilation, and passes it along to the Query Optimizer so that they can be used to generate potentially more efficient query execution plans.

Parameter values are sniffed during compilation or recompilation for the following types of batches:

- Stored procedures
- Queries submitted via `sp_executesql`
- Prepared queries

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-profiling-infrastructure>

QUESTION 6**HOTSPOT**

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

You need to enable automatic tuning for Db1.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: SET AUTOMATIC_TUNING = AUTO To enable automatic tuning on a single database via T-SQL, connect to the database and execute the following query:

ALTER DATABASE current SET AUTOMATIC_TUNING = AUTO

Setting automatic tuning to AUTO will apply Azure Defaults.

Box 2: SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON) To configure individual automatic tuning options via T-SQL, connect to the database and execute the query such as this one:

ALTER DATABASE current SET AUTOMATIC_TUNING (FORCE_LAST_GOOD_PLAN = ON)

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

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-enable>

QUESTION 7 You deploy a database to an Azure SQL Database managed instance.

You need to prevent read queries from blocking queries that are trying to write to the database.

Which database option should set?

- A. PARAMETERIZATION to FORCED
- B. PARAMETERIZATION to SIMPLE
- C. Delayed Durability to **Forced**
- D. READ_COMMITTED_SNAPSHOT to ON

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

In SQL Server, you can also minimize locking contention while protecting transactions from dirty reads of uncommitted data modifications using either:

- The READ COMMITTED isolation level with the READ_COMMITTED_SNAPSHOT database option set to ON. ▪

The SNAPSHOT isolation level.

If READ_COMMITTED_SNAPSHOT is set to ON (the default on SQL Azure Database), the Database Engine uses row versioning to present each statement with a transactionally consistent snapshot of the data as it existed at the start of the statement. Locks are not used to protect the data from updates by other transactions.

Incorrect Answers:

A: When the PARAMETERIZATION database option is set to SIMPLE, the SQL Server query optimizer may choose to parameterize the queries. This means that any literal values that are contained in a query are substituted with parameters. This process is referred to as simple parameterization. When SIMPLE parameterization is in effect, you cannot control which queries are parameterized and which queries are not.

B: You can specify that all queries in a database be parameterized by setting the PARAMETERIZATION database option to FORCED. This process is referred to as forced parameterization.

C: Delayed transaction durability is accomplished using asynchronous log writes to disk. Transaction log records are kept in a buffer and written to disk when the buffer fills or a buffer flushing event takes place. Delayed transaction durability reduces both latency and contention within the system.

Some of the cases in which you could benefit from using delayed transaction durability are:

- You can tolerate some data loss.
- You are experiencing a bottleneck on transaction log writes.
- Your workloads have a high contention rate.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-transaction-isolation-level-transact-sql>

QUESTION 8

You have an Azure SQL database.

You discover that the plan cache is full of compiled plans that were used only once.

You run the `select * from sys.database_scoped_configurations` Transact-SQL command and receive the results shown in the following table.

configuration_id	name	value	is_value_default
1	LEGACY_CARDINALITY_ESTIMATION	0	1
2	QUERY_OPTIMIZER_HOTFIXES	0	1
3	OPTIMIZE_FOR_AD_HOC_WORKLOADS	0	1
4	ACCELERATED_PLAN_FORCING	1	1

You need relieve the memory pressure.

What should you configure?

- A. LEGACY_CARDINALITY_ESTIMATION
- B. QUERY_OPTIMIZER_HOTFIXES
- C. OPTIMIZE_FOR_AD_HOC_WORKLOADS
- D. ACCELERATED_PLAN_FORCING

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

OPTIMIZE_FOR_AD_HOC_WORKLOADS = { ON | OFF }

Enables or disables a compiled plan stub to be stored in cache when a batch is compiled for the first time. The default is OFF. Once the database scoped configuration OPTIMIZE_FOR_AD_HOC_WORKLOADS is enabled for a database, a compiled plan stub will be stored in cache when a batch is compiled for the first time. Plan stubs have a smaller memory footprint compared to the size of the full compiled plan.

Incorrect Answers:

A: LEGACY_CARDINALITY_ESTIMATION = { ON | OFF | PRIMARY }

Enables you to set the query optimizer cardinality estimation model to the SQL Server 2012 and earlier version independent of the compatibility level of the database. The default is OFF, which sets the query optimizer cardinality estimation model based on the compatibility level of the database.

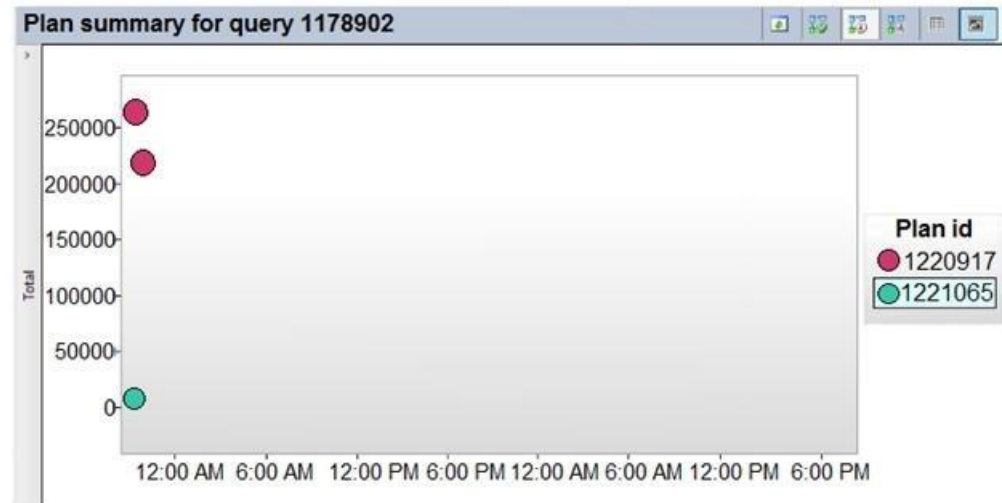
B: QUERY_OPTIMIZER_HOTFIXES = { ON | OFF | PRIMARY }

Enables or disables query optimization hotfixes regardless of the compatibility level of the database. The default is OFF, which disables query optimization hotfixes that were released after the highest available compatibility level was introduced for a specific version (post-RTM).

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/statements/alter-database-scoped-configuration-transact-sql>

QUESTION 9 You have SQL Server on an Azure virtual machine that contains a database named DB1.

You view a plan summary that shows the duration in milliseconds of each execution of query 1178902 as shown in the following exhibit:



What should you do to ensure that the query uses the execution plan which executes in the least amount of time?

- A. Force the query execution plan for plan 1221065.
- B. Run the `DBCC FREEPROCCACHE` command.
- C. Force the query execution plan for plan 1220917.
- D. Disable parameter sniffing.

VCEUp

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:


Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios>

QUESTION 10

HOTSPOT

You have an Azure SQL database named DB1. The automatic tuning options for DB1 are configured as shown in the following exhibit.


 Azure SQL Database built-in intelligence automatically tunes your databases to optimize performance. Click here to learn more about automatic tuning. 


Inherit from: 




Server

Azure defaults

Don't inherit

 The database is inheriting automatic tuning configuration from Azure defaults.

Configure the automatic tuning options 

OPTION	DESIRED STATE	CURRENT STATE
 FORCE PLAN	<div>ON</div> <div>OFF</div> <div>INHERIT</div>	ON Auto-configured by Azure
 CREATE INDEX	<div>ON</div> <div>OFF</div> <div>INHERIT</div>	ON Auto-configured by Azure
 DROP INDEX	<div>ON</div> <div>OFF</div> <div>INHERIT</div>	ON Forced by user

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Nonclustered indexes will be added to tables to improve performance.	<input type="radio"/>	<input type="radio"/>
Columns will be added to existing indexes automatically.	<input type="radio"/>	<input type="radio"/>
The query execution plan will revert to a previous plan if query performance degrades.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
Nonclustered indexes will be added to tables to improve performance.	<input checked="" type="radio"/>	<input type="radio"/>
Columns will be added to existing indexes automatically.	<input type="radio"/>	<input checked="" type="radio"/>
The query execution plan will revert to a previous plan if query performance degrades.	<input checked="" type="radio"/>	<input type="radio"/>

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: Yes

We see: Tuning option: Create index ON

CREATE INDEX - Identifies indexes that may improve performance of your workload, creates indexes, and automatically verifies that performance of queries has improved.

Box 2: No

Box 3: Yes

FORCE LAST GOOD PLAN (automatic plan correction) - Identifies Azure SQL queries using an execution plan that is slower than the previous good plan, and queries using the last known good plan instead of the regressed plan.

QUESTION 11 You have an Azure SQL database named DB1. You run a query while connected to DB1.

You review the actual execution plan for the query, and you add an index to a table referenced by the query. You

need to compare the previous actual execution plan for the query to the Live Query Statistics.

What should you do first in Microsoft SQL Server Management Studio (SSMS)?

- A. For DB1, set QUERY_CAPTURE_MODE of Query Store to All.
- B. Run the `SET SHOWPLAN_ALL` Transact-SQL statement.
- C. Save the actual execution plan.
- D. Enable Query Store for DB1.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The Plan Comparison menu option allows side-by-side comparison of two different execution plans, for easier identification of similarities and changes that explain the different behaviors for all the reasons stated above. This option can compare between:

Two previously saved execution plan files (.sqlplan extension).

One active execution plan and one previously saved query execution plan. Two selected query plans in Query Store.

QUESTION 12

You have an Azure SQL database.

Users report that the executions of a stored procedure are slower than usual. You suspect that a regressed query is causing the performance issue.

You need to view the query execution plan to verify whether a regressed query is causing the issue. The solution must minimize effort.

What should you use?

- A. Performance Recommendations in the Azure portal
- B. Extended Events in Microsoft SQL Server Management Studio (SSMS)
- C. Query Store in Microsoft SQL Server Management Studio (SSMS)
- D. Query Performance Insight in the Azure portal

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Use the Query Store Page in SQL Server Management Studio.

Query performance regressions caused by execution plan changes can be non-trivial and time consuming to resolve.

Since the Query Store retains multiple execution plans per query, it can enforce policies to direct the Query Processor to use a specific execution plan for a query. This is referred to as plan forcing. Plan forcing in Query Store is provided by using a mechanism similar to the USE PLAN query hint, but it does not require any change in user applications. Plan forcing can resolve a query performance regression caused by a plan change in a very short period of time.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store>

QUESTION 13

You have an Azure SQL database. The database contains a table that uses a columnstore index and is accessed infrequently.

You enable columnstore archival compression.

What are two possible results of the configuration? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Queries that use the index will consume more disk I/O.
- B. Queries that use the index will retrieve fewer data pages.
- C. The index will consume more disk space.
- D. The index will consume more memory.
- E. Queries that use the index will consume more CPU resources.

Correct Answer: BE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

For rowstore tables and indexes, use the data compression feature to help reduce the size of the database. In addition to saving space, data compression can help improve performance of I/O intensive workloads because the data is stored in fewer pages and queries need to read fewer pages from disk.

Use columnstore archival compression to further reduce the data size for situations when you can afford extra time and CPU resources to store and retrieve the data.

Testlet 1

Case study

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1, SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1.

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

Requirements

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints. **Business**

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
 - Minimize downtime during the migration of the SERVER1 databases.
 - Use the Azure Hybrid Use Benefits when migrating workloads to Azure. ▪
- Once all requirements are met, minimize costs whenever possible.



QUESTION 1

DRAG DROP

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

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

Select and Place:

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

Correct Answer:

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

Section: [none]
Explanation

Explanation/Reference:
Explanation:

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

1. Create Azure automation account (Step 1)
2. Import SQLServer module (Step 2)
3. Add Credentials to access SQL DB
This will use secure way to hold login name and password that will be used to access Azure SQL DB
4. Add a runbook to run the maintenance (Step 3) Steps:
 1. Click on "runbooks" at the left panel and then click "add a runbook"
 2. Choose "create a new runbook" and then give it a name and choose "Powershell" as the type of the runbook and then click on "create"

The screenshot shows the 'Add Runbook' dialog box. On the left, under 'Quick Create', the 'Create a new runbook' option is selected. On the right, the 'Runbook' form is displayed with the following fields:

- Name:** SqlMaintenance (with a green checkmark)
- Runbook type:** PowerShell (selected from a dropdown menu)
- Description:** (empty text area with a green checkmark)

5. Schedule task (Step 4) Steps:
 1. Click on Schedules
 2. Click on "Add a schedule" and follow the instructions to choose existing schedule or create a new schedule.

Reference: <https://techcommunity.microsoft.com/t5/azure-database-support-blog/automating-azure-sql-db-index-and-statistics-maintenance-using/ba-p/368974>

Question Set 2

QUESTION 1

DRAG DROP

You have SQL Server on an Azure virtual machine named SQL1.

SQL1 has an agent job to back up all databases.

You add a user named dbadmin1 as a SQL Server Agent operator.

You need to ensure that dbadmin1 receives an email alert if a job fails.

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:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Step 1: Enable the email settings for the SQL Server Agent.

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Step 2: Create a job alert

Step 3: Create a job notification Example:

```
-- adds an e-mail notification for the specified alert (Test Alert) -  
- This example assumes that Test Alert already exists -- and  
that François Ajenstat is a valid operator name.  
USE msdb ;  
GO
```

```
EXEC dbo.sp_add_notification  
@alert_name = N'Test Alert',  
@operator_name = N'François Ajenstat',  
@notification_method = 1 ;  
GO
```

Reference:

<https://docs.microsoft.com/en-us/sql/ssms/agent/notify-an-operator-of-job-status>

<https://docs.microsoft.com/en-us/sql/ssms/agent/assign-alerts-to-an-operator>

QUESTION 2

DRAG DROP

You need to apply 20 built-in Azure Policy definitions to all new and existing Azure SQL Database deployments in an Azure subscription. The solution must minimize administrative effort.

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
Duplicate Azure Policy definitions	
Run Azure Policy remediation tasks	
Create an Azure Blueprints assignment	
Create an Azure Policy initiative	
Create an Azure Policy initiative assignment	

Correct Answer:

Actions	Answer Area
Duplicate Azure Policy definitions	Create an Azure Policy initiative
	Create an Azure Policy initiative assignment
Create an Azure Blueprints assignment	Run Azure Policy remediation tasks

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Step 1: Create an Azure Policy Initiative

The first step in enforcing compliance with Azure Policy is to assign a policy definition. A policy definition defines under what condition a policy is enforced and what effect to take.

With an initiative definition, you can group several policy definitions to achieve one overarching goal. An initiative evaluates resources within scope of the assignment for compliance to the included policies.

Step 2: Create an Azure Policy Initiative assignment

Assign the initiative definition you created in the previous step.

Step 3: Run Azure Policy remediation tasks

To apply the Policy Initiative to the existing SQL databases.

Reference: <https://docs.microsoft.com/en-us/azure/governance/policy/tutorials/create-and-manage>

QUESTION 3

You have an Azure SQL Database managed instance named SQLMI1. A Microsoft SQL Server Agent job runs on SQLMI1.

You need to ensure that an automatic email notification is sent once the job completes.

What should you include in the solution?

- A. From SQL Server Configuration Manager (SSMS), enable SQL Server Agent
- B. From SQL Server Management Studio (SSMS), run `sp_set_sqlagent_properties`
- C. From SQL Server Management Studio (SSMS), create a Database Mail profile
- D. From the Azure portal, create an Azure Monitor action group that has an Email/SMS/Push/Voice action

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

To send a notification in response to an alert, you must first configure SQL Server Agent to send mail.

Using SQL Server Management Studio; to configure SQL Server Agent to use Database Mail:

1. In Object Explorer, expand a SQL Server instance.
2. Right-click SQL Server Agent, and then click Properties.
3. Click Alert System.
4. Select Enable Mail Profile.
5. In the Mail system list, select Database Mail.
6. In the Mail profile list, select a mail profile for Database Mail.
7. Restart SQL Server Agent.

Note: Prerequisites include:

- Enable Database Mail.
- Create a Database Mail account for the SQL Server Agent service account to use.
 - Create a Database Mail profile for the SQL Server Agent service account to use and add the user to the DatabaseMailUserRole in the msdb database.
- Set the profile as the default profile for the msdb database.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/database-mail/configure-sql-server-agent-mail-to-use-database-mail>

Testlet 1

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.

- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data ▪
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints. **Business**

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
 - Minimize downtime during the migration of the SERVER1 databases.
 - Use the Azure Hybrid Use Benefits when migrating workloads to Azure. ▪
- Once all requirements are met, minimize costs whenever possible.

QUESTION 1 HOTSPOT

You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure. The solution must meet the business requirements.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Quorum mode:

	▼
Node majority with witness	
Node majority	
No majority	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	
Azure Standard Load Balancer	

Correct Answer:**Answer Area**

Quorum mode:

	▼
Node majority with witness	
Node majority	
No majority	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	
Azure Standard Load Balancer	

Section: [none]**Explanation****Explanation/Reference:**

Explanation:

Box 1: Node majority with witness

As a general rule when you configure a quorum, the voting elements in the cluster should be an odd number. Therefore, if the cluster contains an even number of voting nodes, you should configure a disk witness or a file share witness.

Note: Mode: Node majority with witness (disk or file share)

Nodes have votes. In addition, a quorum witness has a vote. The cluster quorum is the majority of voting nodes in the active cluster membership plus a witness vote. A quorum witness can be a designated disk witness or a designated file share witness.

Box 2: Azure Standard Load Balancer

Microsoft guarantees that a Load Balanced Endpoint using Azure Standard Load Balancer, serving two or more Healthy Virtual Machine Instances, will be available 99.99% of the time.

Scenario: Business Requirements

Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.

Incorrect Aswers:

Basic Balancer: No SLA is provided for Basic Load Balancer.

Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a public Virtual IP (VIP) that is accessible over the internet.

An internal listener uses an internal load balancer and only supports clients within the same Virtual Network.

Reference:

<https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness>

https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/

QUESTION 2 What should you do after a failover of SalesSQLDb1 to ensure that the database remains accessible to SalesSQLDb1App1?

- A. Configure SalesSQLDb1 as writable.
- B. Update the connection strings of SalesSQLDb1App1.
- C. Update the firewall rules of SalesSQLDb1.
- D. Update the users in SalesSQLDb1.

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Scenario: SalesSQLDb1 uses database firewall rules and contained database users.

Testlet 2

Case study

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

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

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

To start the case study

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

Overview

Existing Environment

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

Active Directory

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

Database Environment

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

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

VCEUp

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

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

Current Business Model

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

Requirements

Planned Changes

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

The estimated size of each PaaS database is 1 TB.

Contoso plans to implement the following changes:

- Move the PostgreSQL database to Azure Database for PostgreSQL during the next six months.
 - Upgrade SQL1, SQL2, and SQL3 to SQL Server 2019 during the next few months. ▪
- Start onboarding customers to the new PaaS solution within six months.

Business Goals

Contoso identifies the following business requirements:

- Use built-in Azure features whenever possible.
- Minimize development effort whenever possible.
- Minimize the compute costs of the PaaS solutions.
- Provide all the customers with their own copy of the database by using the PaaS solution.
- Provide the customers with different table and row access based on the customer's service agreement.
- In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.
- Ensure that users of the PaaS solution can create their own database objects but be prevented from modifying any of the existing database objects supplied by Contoso. **Technical**

Requirements

Contoso identifies the following technical requirements:

- Users of the PaaS solution must be able to sign in by using their own corporate Azure AD credentials or have Azure AD credentials supplied to them by Contoso. The solution must avoid using the internal Azure AD of Contoso to minimize guest users.
- All customers must have their own resource group, Azure SQL server, and Azure SQL database. The deployment of resources for each customer must be done in a consistent fashion.
- Users must be able to review the queries issued against the PaaS databases and identify any new objects created.
- Downtime during the PostgreSQL database migration must be minimized.

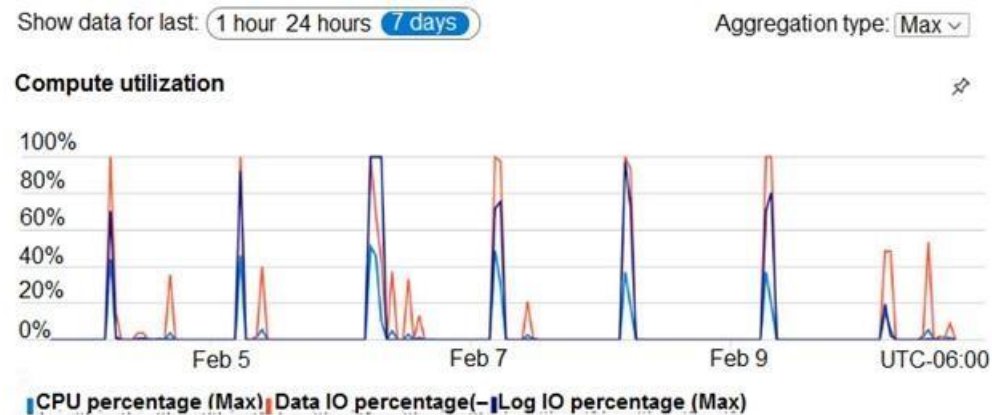
Monitoring Requirements

Contoso identifies the following monitoring requirements:

- Notify administrators when a PaaS database has a higher than average CPU usage.
- Use a single dashboard to review security and audit data for all the PaaS databases.
- Use a single dashboard to monitor query performance and bottlenecks across all the PaaS databases.
- Monitor the PaaS databases to identify poorly performing queries and resolve query performance issues automatically whenever possible.

PaaS Prototype

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



Role Assignments

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

[+ Add](#)
[Edit columns](#)
[Refresh](#)
[Remove](#)
[Got feedback?](#)

[Check access](#)
[Role assignments](#)
[Deny assignments](#)
[Classic administrators](#)
[Roles](#)

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

Number of role assignments for this subscription ⓘ



15 2000

Name ⓘ
 Type ⓘ
 Role ⓘ
 Scope ⓘ

Group by ⓘ

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

2 items (2 Groups)

<input type="checkbox"/>	Name	Type	Role	Scope
Contributor				
<input type="checkbox"/>	 DBAGroup1	Group	Contributor ⓘ	This resource
SQL DB Contributor				
<input type="checkbox"/>	 DBAGroup2	Group	SQL DB Contributor ⓘ	This resource

QUESTION 1 What should you implement to meet the disaster recovery requirements for the PaaS solution?

- A. Availability Zones
- B. failover groups
- C. Always On availability groups
- D. geo-replication

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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

The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active georeplication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.

The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

Question Set 3

QUESTION 1

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: From the Azure portal, you delete Database1 from Server2, and then you create a new database on Server2 by using the backup of Database1 from Server1.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

QUESTION 2

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: You run the `Remove-AzSqlDatabase` PowerShell cmdlet for Database1 on Server2. You run the `Restore-AzSqlDatabase` PowerShell cmdlet for Database1 on Server2.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

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

QUESTION 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: You restore Database1 from Server1 to the Server2 by using the `RESTORE Transact-SQL` command and the `REPLACE` option.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The `REPLACE` option overrides several important safety checks that restore normally performs. The overridden checks are as follows: •

Restoring over an existing database with a backup taken of another database.

With the `REPLACE` option, restore allows you to overwrite an existing database with whatever database is in the backup set, even if the specified database name differs from the database name recorded in the backup set. This can result in accidentally overwriting a database by a different database.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-transact-sql>

QUESTION 4 You have an Always On availability group deployed to Azure virtual machines. The availability group contains a database named DB1 and has two nodes named SQL1 and SQL2. SQL1 is the primary replica.

You need to initiate a full backup of DB1 on SQL2.

Which statement should you run?

- A. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (Differential, STATS=5, COMPRESSION);`
- B. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (COPY_ONLY, STATS=5, COMPRESSION);`
- C. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (File_Snapshot, STATS=5, COMPRESSION);`
- D. `BACKUP DATABASE DB1 TO URL='https://mystorageaccount.blob.core.windows.net/mycontainer/DB1.bak' with (NoInit, STATS=5, COMPRESSION);`

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

`BACKUP DATABASE` supports only copy-only full backups of databases, files, or filegroups when it's executed on secondary replicas. Copy-only backups don't impact the log chain or clear the differential bitmap.

Incorrect Answers:

A: Differential backups are not supported on secondary replicas. The software displays this error because the secondary replicas support copy-only database backups.

Reference: <https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/active-secondaries-backup-on-secondary-replicas-always-on-availability-groups>

QUESTION 5

You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

- A. An Azure SQL Database elastic pool
- B. SQL Server on Azure virtual machines
- C. an Azure SQL Database managed instance
- D. Azure SQL databases

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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

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

QUESTION 6 You have 10 Azure virtual machines that have SQL Server installed.

You need to implement a backup strategy to ensure that you can restore specific databases to other SQL Server instances. The solution must provide centralized management of the backups.

What should you include in the backup strategy?

- A. Automated Backup in the SQL virtual machine settings
- B. Azure Backup
- C. Azure Site Recovery
- D. SQL Server Agent jobs

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure Backup provides an Enterprise class backup capability for SQL Server on Azure VMs. All backups are stored and managed in a Recovery Services vault. There are several advantages that this solution provides, especially for Enterprises.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/backup-restore#azbackup>

QUESTION 7

You need to recommend an availability strategy for an Azure SQL database. The strategy must meet the following requirements:

- Support failovers that do not require client applications to change their connection strings.
 - Replicate the database to a secondary Azure region. ▪
- Support failover to the secondary region.

What should you include in the recommendation?

- A. failover groups
- B. transactional replication
- C. Availability Zones
- D. geo-replication

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Active geo-replication is an Azure SQL Database feature that allows you to create readable secondary databases of individual databases on a server in the same or different data center (region).

Incorrect Answers:

C: Availability Zones are unique physical locations within a region. Each zone is made up of one or more datacenters equipped with independent power, cooling, and networking.

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview>

QUESTION 8

DRAG DROP

You have SQL Server on an Azure virtual machine that contains a database named DB1. DB1 is 30 TB and has a 1-GB daily rate of change.

You back up the database by using a Microsoft SQL Server Agent job that runs Transact-SQL commands. You perform a weekly full backup on Sunday, daily differential backups at 01:00, and transaction log backups every five minutes.

The database fails on Wednesday at 10:00.

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

Select and Place:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

VCEUp

QUESTION 9

You are building a database backup solution for a SQL Server database hosted on an Azure virtual machine.

In the event of an Azure regional outage, you need to be able to restore the database backups. The solution must minimize costs.

Which type of storage accounts should you use for the backups?

- A. locally-redundant storage (LRS)
- B. read-access geo-redundant storage (RA-GRS)
- C. zone-redundant storage (ZRS)
- D. geo-redundant storage

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Geo-redundant storage (with GRS or GZRS) replicates your data to another physical location in the secondary region to protect against regional outages. However, that data is available to be read only if the customer or Microsoft initiates a failover from the primary to secondary region. When you enable read access to the secondary region, your data is available to be read if the primary region becomes unavailable. For read access to the secondary region, enable read-access geo-redundant storage (RA-GRS) or read-access geo-zone-redundant storage (RA-GZRS).

Incorrect Answers:

A: Locally redundant storage (LRS) copies your data synchronously three times within a single physical location in the primary region. LRS is the least expensive replication option, but is not recommended for applications requiring high availability.

C: Zone-redundant storage (ZRS) copies your data synchronously across three Azure availability zones in the primary region.

D: Geo-redundant storage (with GRS or GZRS) replicates your data to another physical location in the secondary region to protect against regional outages. However, that data is available to be read only if the customer or Microsoft initiates a failover from the primary to secondary region.

Reference: <https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>

QUESTION 10 You have SQL Server on Azure virtual machines in an availability group.

You have a database named DB1 that is **NOT** in the availability group.

You create a full database backup of DB1.

You need to add DB1 to the availability group.

Which restore option should you use on the secondary replica?

- A. Restore with Recovery
- B. Restore with Norecovery
- C. Restore with Standby

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Prepare a secondary database for an Always On availability group requires two steps:

1. Restore a recent database backup of the primary database and subsequent log backups onto each server instance that hosts the secondary replica, using RESTORE WITH NORECOVERY
2. Join the restored database to the availability group.

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server>

QUESTION 11

You are planning disaster recovery for the failover group of an Azure SQL Database managed instance.

Your company's SLA requires that the database in the failover group become available as quickly as possible if a major outage occurs.

You set the Read/Write failover policy to **Automatic**.

What are two results of the configuration? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. In the event of a datacenter or Azure regional outage, the databases will fail over automatically.
- B. In the event of an outage, the databases in the primary instance will fail over immediately.
- C. In the event of an outage, you can selectively fail over individual databases.
- D. In the event of an outage, you can set a different grace period to fail over each database.
- E. In the event of an outage, the minimum delay for the databases to fail over in the primary instance will be one hour.

Correct Answer: AE

Section: [none]

Explanation

Explanation/Reference:

Explanation:

A: Auto-failover groups allow you to manage replication and failover of a group of databases on a server or all databases in a managed instance to another region.

E: Because verification of the scale of the outage and how quickly it can be mitigated involves human actions by the operations team, the grace period cannot be set below one hour. This limitation applies to all databases in the failover group regardless of their data synchronization state.

Incorrect Answers:

C: individual SQL Managed Instance databases cannot be added to or removed from a failover group.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

QUESTION 12

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Premium
- B. Azure SQL Database serverless
- C. Azure SQL Database managed instance Business Critical
- D. Azure SQL Database Standard

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW). The routing to a specific gateway ring is controlled by Azure Traffic Manager (ATM). Because the zone redundant configuration in the Premium or Business Critical service tiers does not create additional database redundancy, you can enable it at no extra cost. By selecting a zone redundant configuration, you can make your Premium or Business Critical databases resilient to a much larger set of failures, including catastrophic datacenter outages, without any changes to the application logic. You can also convert any existing Premium or Business Critical databases or pools to the zone redundant configuration.

Incorrect Answers:

C. This feature is not available in SQL Managed Instance.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

QUESTION 13

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

- During normal operations, provide at least two readable copies of Sales. ▪
- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and geo-replication.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation**Explanation/Reference:**

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

QUESTION 14

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

- During normal operations, provide at least two readable copies of Sales.
- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

QUESTION 15

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure SQL database named Sales.

You need to implement disaster recovery for Sales to meet the following requirements:

- During normal operations, provide at least two readable copies of Sales. ▪
- Ensure that Sales remains available if a datacenter fails.

Solution: You deploy an Azure SQL database that uses the General Purpose service tier and failover groups.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Instead deploy an Azure SQL database that uses the Business Critical service tier and Availability Zones.

Note: Premium and Business Critical service tiers leverage the Premium availability model, which integrates compute resources (sqlservr.exe process) and storage (locally attached SSD) on a single node. High availability is achieved by replicating both compute and storage to additional nodes creating a three to four-node cluster.

By default, the cluster of nodes for the premium availability model is created in the same datacenter. With the introduction of Azure Availability Zones, SQL Database can place different replicas of the Business Critical database to different availability zones in the same region. To eliminate a single point of failure, the control ring is also duplicated across multiple zones as three gateway rings (GW).

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/high-availability-sla>

QUESTION 16 You plan to move two 100-GB databases to Azure.

You need to dynamically scale resources consumption based on workloads. The solution must minimize downtime during scaling operations.

What should you use?

- A. two Azure SQL Databases in an elastic pool
- B. two databases hosted in SQL Server on an Azure virtual machine
- C. two databases in an Azure SQL Managed instance
- D. two single Azure SQL databases

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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

Reference:

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

QUESTION 17

You have an on-premises app named App1 that stores data in an on-premises Microsoft SQL Server 2016 database named DB1.

You plan to deploy additional instances of App1 to separate Azure regions. Each region will have a separate instance of App1 and DB1. The separate instances of DB1 will sync by using Azure SQL Data Sync.

You need to recommend a database service for the deployment. The solution must minimize administrative effort.

What should you include in the recommendation?

- A. Azure SQL Managed instance

- B. Azure SQL Database single database
- C. Azure Database for PostgreSQL
- D. SQL Server on Azure virtual machines

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure SQL Database single database supports Data Sync.

Incorrect Answers:

A: Azure SQL Managed instance does not support Data Sync.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison>

QUESTION 18

You have an Azure SQL database named DB1.

You need to ensure that DB1 will support automatic failover without data loss if a datacenter fails. The solution must minimize costs.

Which deployment option and pricing tier should you configure?

- A. Azure SQL Database Premium
- B. Azure SQL Database Hyperscale
- C. Azure SQL Database managed instance Business Critical
- D. Azure SQL Database Standard

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Auto-failover groups is required. SQL Managed Instance supports Auto-failover groups.

Note: The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region.

Incorrect Answers:

B: Hyperscale is for large databases. It is designed for most business workloads, providing highly scalable storage, read scale-out, and fast database restore capabilities.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

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 stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have two Azure SQL Database servers named Server1 and Server2. Each server contains an Azure SQL database named Database1.

You need to restore Database1 from Server1 to Server2. The solution must replace the existing Database1 on Server2.

Solution: From Microsoft SQL Server Management Studio (SSMS), you rename Database1 on Server2 as Database2. From the Azure portal, you create a new database on Server2 by restoring the backup of Database1 from Server1, and then you delete Database2.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Instead restore Database1 from Server1 to the Server2 by using the RESTORE Transact-SQL command and the REPLACE option.

Note: REPLACE should be used rarely and only after careful consideration. Restore normally prevents accidentally overwriting a database with a different database. If the database specified in a RESTORE statement already exists on the current server and the specified database family GUID differs from the database family GUID recorded in the backup set, the database is not restored. This is an important safeguard.

Reference:

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

Testlet 1

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

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

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

To start the case study

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

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

Identity Environment

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant. **Database**

Environment

The sales department has the following database workload:

- An on-premises server named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contains database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1.

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

Licensing Agreement

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

Current Problems

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

Requirements

Planned Changes

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.
- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

Technical Requirements

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Security and Compliance Requirements

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints. **Business**

Requirements

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
 - Minimize downtime during the migration of the SERVER1 databases.
 - Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- Once all requirements are met, minimize costs whenever possible.

QUESTION 1 You need to implement authentication for ResearchDB1. The solution must meet the security and compliance requirements.

What should you run as part of the implementation?

- A. CREATE LOGIN and the FROM WINDOWS clause
- B. CREATE USER and the FROM CERTIFICATE clause
- C. CREATE USER and the FROM LOGIN clause
- D. CREATE USER and the ASYMMETRIC KEY clause
- E. CREATE USER and the FROM EXTERNAL PROVIDER clause

Correct Answer: E

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.)

Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user:

```
CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;
```

Reference:

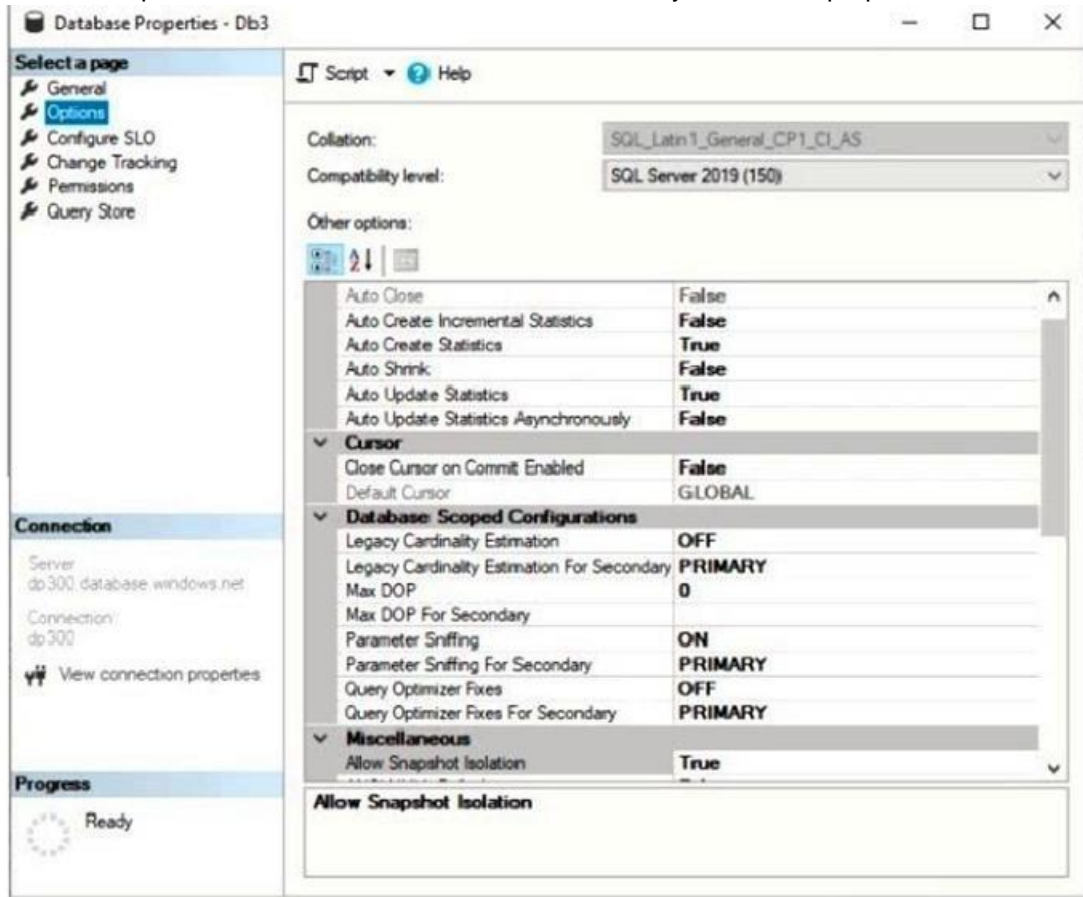
<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql>

Question Set 2

QUESTION 1

You have an Azure SQL database named DB3.

You need to provide a user named DevUser with the ability to view the properties of DB3 from Microsoft SQL Server Management Studio (SSMS) as shown in the exhibit. (Click the **Exhibit** tab.)



VCEûp

Which Transact-SQL command should you run?

- A. GRANT SHOWPLAN TO DevUser
- B. GRANT VIEW DEFINITION TO DevUser
- C. GRANT VIEW DATABASE STATE TO DevUser
- D. GRANT SELECT TO DevUser

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The exhibit displays Database [State] properties.

To query a dynamic management view or function requires SELECT permission on object and VIEW SERVER STATE or VIEW DATABASE STATE permission.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/database-properties-options-page>

QUESTION 2

HOTSPOT

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

The database reports a CHECKSUM error.

You need to recover the database.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: SINGLE_USER The specified database must be in single-user mode to use one of the following repair options.

Box 2: REPAIR_ALLOW_DATA_LOSS REPAIR_ALLOW_DATA_LOSS tries to repair all reported errors. These repairs can cause some data loss.

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

Incorrect Answers:

REPAIR_FAST Maintains syntax for backward compatibility only. No repair actions are performed.

Box 3: MULTI_USER MULTI_USER All users that have the appropriate permissions to connect to the database are allowed.

Reference: <https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql>

QUESTION 3

HOTSPOT

You have an Azure SQL Database managed instance named sqldbmi1 that contains a database name Sales.

You need to initiate a backup of Sales.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Correct Answer:

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: TO URL = 'https://storage1.blob.core.windows.net/blob1/Sales.bak'
Native database backup in Azure SQL Managed Instance.
You can backup any database using standard BACKUP T-SQL command:

BACKUP DATABASE tpcc2501

TO URL = 'https://myacc.blob.core.windows.net/testcontainer/tpcc2501.bak'

WITH COPY_ONLY Box 2: WITH COPY_ONLY

Reference: <https://techcommunity.microsoft.com/t5/azure-sql-database/native-database-backup-in-azure-sql-managed-instance/ba-p/386154>

QUESTION 4

You have the following Transact-SQL query.

```
SELECT
    [file_id] AS [File ID],
    [type] AS [File Type],
    substring([physical_name], 1,1) AS [Drive],
    [name] AS [Logical Name],
    [physical_name] AS [Physical Name],
    CAST([size] as DECIMAL(38,0))/128.0 AS [ColumnA],
    CAST(FILEPROPERTY([name], 'SpaceUsed') AS DECIMAL(38,0))/128.0 AS
[ColumnB],
    (CAST([size] AS DECIMAL(38,0))/128.0) - (CAST(FILEPROPERTY([name],
'SpaceUsed') AS DECIMAL (38,0))/128.0) AS [ColumnC],
    [max_size] AS [ColumnD],
    [is_percent_growth] AS [Percent Growth Enabled],
    [growth] AS [Growth Rate],
    SYSDATETIME() AS [Current Date]
FROM sys.database_files;
```

Which column returned by the query represents the free space in each file?

- A. ColumnA
- B. ColumnB
- C. ColumnC
- D. ColumnD

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Example:

Free space for the file in the below query result set will be returned by the FreeSpaceMB column.

```
SELECT DB_NAME() AS DbName,
name AS FileName, type_desc,
size/128.0 AS CurrentSizeMB,
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpaceMB
FROM sys.database_files
WHERE type IN (0,1);
```

Reference: <https://www.sqlshack.com/how-to-determine-free-space-and-file-size-for-sql-server-databases/>