DP-300

Number: DP-300
Passing Score: 800
Time Limit: 120 min
File Version: 1

DP-300

https://www.gratisexam.com/

Plan and Implement Data Platform Resources

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.

Existing Environment

Network Environment

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

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.

Plan and Implement Data Platform Resources

Question Set 2

QUESTION 1
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.

GRATIS EXAM
Free Practice Exams
https://www.gratisexam.com/

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 2
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 3
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: MAX(<Total number of DBs X average vCore utilization per DB>, <Number of concurrently peaking DBs X 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

Implement a Secure Environment

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.

Existing Environment

Network Environment

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

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.

Implement a Secure Environment

Question Set 2

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
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	1	847	847	C244A066-2271-...	Confidential - GDPR
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	2	32	32	C244A066-2271-...	Confidential
	[redacted]7.125	Microsoft SQL Server Management Studio - Query	41	32	32	A7088FD4-759E-...	Confidential, Confidential - GDPR

https://www.gratisexam.com/

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

Monitor and Optimize Operational Resources

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.

Existing Environment

Network Environment

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

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.

Monitor and Optimize Operational Resources

Question Set 2

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 VM1
D. 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 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
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

https://www.gratisexam.com/

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 7
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 vCPUSs 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

Optimize Query Performance

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.

Existing Environment

Network Environment

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

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
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

Optimize Query Performance

Question Set 2

QUESTION 1
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 2
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.

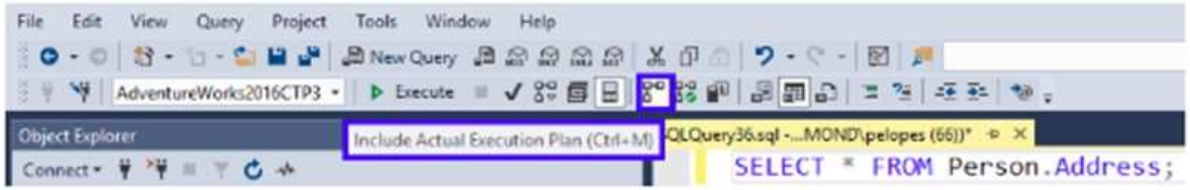

Note: Actual execution plans are generated after the Transact-SQL queries or batches execute. Because of this, an actual execution plan contains runtime information, such as actual resource usage metrics and runtime warnings (if any). The execution plan that is generated displays the actual query execution 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 3
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 4
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 5
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 6
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

Perform Automation of Tasks

Question Set 1

QUESTION 1
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:

GRATIS EXAM
Free Practice Exams
https://www.gratisexam.com/

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

Plan and Implement a High Availability and Disaster Recovery (HADR) Environment

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.

Existing Environment

Network Environment

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

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.

Plan and Implement a High Availability and Disaster Recovery (HADR) Environment

Question Set 2

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 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 3
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 4
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

GRATIS EXAM
Free Practice Exams
https://www.gratisexam.com/