

Microsoft.DP-300J.v2026-08-17.q171

試験コード : DP-300J
試験名称 : Administering Relational Databases on Microsoft Azure (DP-300日本語版)
認証ベンダー : Microsoft
無料問題の数 : 171
バージョン : v2026-08-17
ページの閲覧量 : 102
問題集の閲覧量 : 1838

<https://www.jpnshtken.com/shiken/Microsoft.DP-300J.v2026-08-17.q171.html>

質問: 1

あなたは、Server1という名前の仮想マシン上にSQL Serverインスタンスを含むAzureサブスクリプションを所有しています。
Server1上で以下のT-SQLステートメントを実行します。

```
EXEC msdb.dbo.sp_add_alert @name=N'Fatal Error',
    @message_id=0,
    @severity=20,
    @enabled=1,
    @delay_between_responses=0,
    @include_event_description_in=1;
GO

EXEC msdb.dbo.sp_add_operator
    @name = 'DBA-Team',
    @enabled = 1,
    @email_address = 'DBA-Team@contoso.com';
GO

EXEC msdb.dbo.sp_add_notification
    @alert_name = N'Fatal Error',
    @operator_name = 'DBA-Team',
    @notification_method = 1;
GO
```

以下の各記述について、正しい場合は「はい」を、そうでない場合は「いいえ」を選択してください。
注：正解ごとに1ポイントが加算されます。

Statements	Yes	No
An alert will be generated when a fatal error occurs.	☐	☐
A SQL Server Agent job will run when an alert is generated.	☐	☐
An email will be sent to DBA-Team@contoso.com when an alert is generated.	☐	☐

正解:

Answer Area

Statements	Yes	No
An alert will be generated when a fatal error occurs.	☒	☐
A SQL Server Agent job will run when an alert is generated.	☐	☒
An email will be sent to DBA-Team@contoso.com when an alert is generated.	☒	☐

Explanation:

Answer Area

Statements	Yes	No
An alert will be generated when a fatal error occurs.	☒	☐
A SQL Server Agent job will run when an alert is generated.	☐	☒
An email will be sent to DBA-Team@contoso.com when an alert is generated.	☒	☐

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure subscription that contains a SQL Server instance on a virtual machine named Server1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 2

SQL1 という名前の Azure SQL データベースを含む Azure サブスクリプションがあります。

SQL1 は、可用性ゾーンをサポートしていない Azure リージョンにあります。

同じリージョンに SQL1 のセカンダリ レプリカがあることを確認する必要があります。

何を使うべきですか？

- A. ログ配布
- B. 自動フェイルオーバー グループ
- C. アクティブ geo レプリケーション
- D. Microsoft SQL Server フェールオーバー クラスタ

正解: [\(正解を表示します\)](#)

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: Active geo-replication creates readable secondary databases, but failover groups add the managed listener and automatic failover abstraction needed for application continuity at scale. This fits the case because the requested outcome is: You need to ensure that you have a secondary replica of SQLI in the same region. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Log shipping is a SQL Server HA/DR technique that copies and restores log backups to a secondary database; it is more manual and operationally heavier for Azure SQL Managed Instance migration than LRS. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why B is Wrong: Auto-failover groups provide managed geo-replication and listener-based failover for Azure SQL Database or Managed Instance, including automatic failover policies for regional outages. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: Microsoft SQL Server failover clusters is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: 3

ハイパースケール サービス層に、SQLDB1 と SQLDB2 という名前の Azure SQL データベースが 2 つあります。SQLDB1 は、データベース作成時にハイパースケール サービス層に設定されました。SQLDB2 は、過去 14 日間にビジネス クリティカル サービス層からハイパースケール サービス層に変更されました。

SQLDB1とSQLDB2をビジネスクリティカルサービス層に変更する必要があります。このソリューションは、管理作業を最小限に抑えるものでなければなりません。

まず最初に何をすべきでしょうか？回答欄で適切な選択肢を選んでください。


注：正解ごとに1ポイントが加算されます。

Answer Area



正解:

Answer Area



Explanation:



Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to change SQLDB1 and SQLDB2 to the Business Critical service tier.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 4

Server1 という名前の Azure SQL マネージド インスタンスと、Microsoft SQL Server データベースのバックアップ ファイルが含まれる storage1 という名前の Azure Blob Storage アカウントがあります。

Log Replay Service を使用して、バックアップファイルを storage1 から Server1 に移行する予定です。このソリューションでは、storage1 への接続時に最高レベルのセキュリティを使用する必要があります。

接続を安全にするために、どの PowerShell コマンドレットを実行し、どのパラメーターを指定する必要がありますか？

回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Microsoft

Cmdlet: Start-AzSqlInstanceDatabaseLogReplay
Complete-AzSqlInstanceDatabaseLogReplay
Get-AzSqlInstanceDatabaseLogReplay
Start-AzSqlInstanceDatabaseLogReplay

Parameter: ManagedIdentity
ManagedIdentity
SharedAccessSignature
StorageAccountKey

正解:

ANSWER AREA

Microsoft

Cmdlet: Start-AzSqlInstanceDatabaseLogReplay
Complete-AzSqlInstanceDatabaseLogReplay
Get-AzSqlInstanceDatabaseLogReplay
Start-AzSqlInstanceDatabaseLogReplay

Parameter: ManagedIdentity
ManagedIdentity
SharedAccessSignature
StorageAccountKey

Explanation:

Answer Area

Microsoft

Cmdlet: Start-AzSqlInstanceDatabaseLogReplay

Parameter: ManagedIdentity

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: The solution must use the highest level of security when connecting to storage1.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 5

Microsoft SQL Server がインストールされた Server1 という名前の Azure 仮想マシンがあります。Server1 には DB1 という名前のデータベースが含まれています。

ADB1 という名前の Azure SQL データベースを含む ASVR1 という名前の論理 SQL サーバーがあります。

SQL データ同期を使用して、DB1 を Server1 から ASVR1 に移行する予定です。

移行のための環境を準備する必要があります。ソリューションでは、Server1からADB1への接続がパブリックエンドポイントを使用しないようにする必要があります。

どうすればいいのでしょうか？ 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

Microsoft

Sync data by using:

- Azure Private Link
- A service endpoint
- Azure Private Link
- The SQL Data Sync private link

Set the database isolation level to:

- SNAPSHOT ISOLATION
- REPEATABLE READ
- SERIALIZABLE
- SNAPSHOT ISOLATION

正解:

Answer Area

Microsoft

Sync data by using:

- Azure Private Link
- A service endpoint
- Azure Private Link
- The SQL Data Sync private link

Set the database isolation level to:

- SNAPSHOT ISOLATION
- REPEATABLE READ
- SERIALIZABLE
- SNAPSHOT ISOLATION

Explanation:

Answer Area

Microsoft

Sync data by using:

- Azure Private Link

Set the database isolation level to:

- SNAPSHOT ISOLATION

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to prepare the environment for the migration.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 6

RG1 というリソース グループを含む Azure サブスクリプションがあります。RG1 には、Azure Virtual Machines 上の SQL という SQL Server のインスタンスが含まれています。PowerShell を使用して、SQL の自動修正を有効にして構成する必要があります。ソリューションには、SQL Server と Windows の両方のセキュリティ更新プログラムが含まれている必要があります。

コマンドをどのように完了すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"
```



Microsoft

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
```

Set-AzVMExtension
Get-AzVMSQLServerExtension
Set-AzVMExtension
Set-AzVMSqlServerExtension

```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

```
-SQLManagementType
```

-Lightweight
-Full
-Lightweight
-NoAgent

正解:

Answer Area



Microsoft

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable
```

```
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2
```

```
-MaintenanceWindowDuration 120 -PatchCategory "Important"
```

```
Get-AzVM -ResourceGroupName "RG1" -Name "SQ1" |
```

Set-AzVMExtension
Get-AzVMSQLServerExtension
Set-AzVMExtension
Set-AzVMSqlServerExtension

```
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM
```

```
-SQLManagementType
```

-Lightweight
-Full
-Lightweight
-NoAgent

Explanation:

Answer Area

```
$AutoPatchingConfig = New-AzVMSqlServerAutoPatchingConfig -Enable  
-DayOfWeek "Sunday" -MaintenanceWindowStartingHour 2  
-MaintenanceWindowDuration 120 -PatchCategory "Important"  
  
Get-AzVM -ResourceGroupName "RG1" -Name "SQL1" | Set-AzVMExtension  
-AutoPatchingSettings $AutoPatchingConfig | Update-AzVM  
-SQLManagementType "Lightweight"
```

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: RG1 contains an instance of SQL Server on Azure Virtual Machines named SQL. You need to use PowerShell to enable and configure automated patching for SQL. The solution must include both SQL Server and Windows security updates.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 7

DW1という名前のAzureSynapseAnalyticsのエンタープライズデータウェアハウスを含むAzureデータソリューションがあります。

複数のユーザーがDW1へのアドホッククエリを同時に実行します。

DW1への自動データロードを定期的に行います。

自動データロードに、アドホッククエリの実行時に迅速かつ正常に完了するのに十分なメモリが用意されていることを確認する必要があります。

あなたは何をすべきか？

- A. 自動化されたデータロードクエリに、より小さなリソースクラスを割り当てます。
- B. DW1の各テーブルのすべての列にサンプル統計を作成します。
- C. 自動化されたデータロードクエリにより大きなリソースクラスを割り当てます。
- D. 自動データロードを実行する前に、DW1で大きなファクトテーブルをハッシュ分散します。

正解: ([正解を表示します](#))

The performance capacity of a query is determined by the user's resource class.

Smaller resource classes reduce the maximum memory per query, but increase concurrency.

Larger resource classes increase the maximum memory per query, but reduce concurrency.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/resource-classes-for-workloadmanagement> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why C is Correct: Assign a larger resource class to the automated data load queries. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You have an Azure data solution that contains an enterprise data warehouse in Azure Synapse Analytics named DW1. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Assign a smaller resource class to the automated data load queries. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Create sampled statistics to every column in each table of DW1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: Hash distribute the large fact tables in DW1 before performing the automated data loads. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 8

SQL1という名前のAzure SQLデータベースがあります。

動的管理ビューを使用して、SQL1のリソース使用状況を監視する必要があります。ソリューションは、平均メモリ使用量とピークCPU使用率を返す必要があります。

質問にはどのように回答すればよいですか？回答するには、回答欄で適切なオプションを選択してください。

注：正解ごとに1ポイントが加算されます。

Answer Area

```
SELECT database_name = SQL1(),  
       AVG(avg_memory_usage_percent) AS 'Average memory usage'  
       MAX(  
FROM sys.  
       ) AS 'Maximum CPU usage'
```

avg_cpu_percent
cpu_limit
max_worker_percent

dm_db_resource_stats
dm_db_wait_stats
dm_operation_status

正解:

Answer Area

Microsoft

```

SELECT database_name = SQL1(),
       AVG(avg_memory_usage_percent) AS 'Average memory usage',
       MAX(
         avg_cpu_percent
       ) AS 'Maximum CPU usage'
FROM sys.

```

dm_db_resource_stats

Explanation:

Answer Area

Microsoft

```

SELECT database_name = SQL1(),
       AVG(avg_memory_usage_percent) AS 'Average memory usage',
       MAX(
         max_worker_percent
       ) AS 'Maximum CPU usage'
FROM sys.

```

dm_db_resource_stats

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to monitor the resource usage of SQL1 by using a dynamic management view.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

質問: 9

DB1 という名前の Azure SQL データベースがあります。

User 1 という名前のユーザーには Azure AD アカウントがあります。

User1 に、DBV 内のテーブルに列を追加したり削除したりする権限を与える必要があります。ソリューションでは、最小権限の原則を使用する必要があります。

実行すべき 2 つのアクションはどれですか。それぞれの正解は解決策の一部を示しています。

注: 正解ごとに1ポイント獲得

- A. データベース ユーザーに db.ddladmin ロールを割り当てます。
- B. データベース ユーザーに db.owner ロールを割り当てます。
- C. 包含データベース ユーザーを作成します。
- D. ログインと関連するデータベース ユーザーを作成します。

正解: ([正解を表示します](#))

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why A and D is Correct: Assign the database user the db.ddladmin role. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to provide User1 with the ability to add and remove columns from the tables in DBV The solution must use the principle of least privilege. Assign the database user the db.ddladmin role. satisfies that constraint without adding an unrelated service or manual process. Create a login and an associated database user. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario asks for: You need to provide User1 with the ability to add and remove columns from the tables in DBV The solution must use the principle of least privilege. That makes Create a login and an associated database user. the option that satisfies the required Azure SQL layer and operational outcome.

Why B is Wrong: Assign the database user the db.owner role. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to provide User1 with the ability to add and remove columns from the tables in DBV The solution must use the principle of least privilege.

Why C is Wrong: Create a contained database user. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

質問: 10

server1 というサーバーに db1 という Azure SQL データベースがあります。

db1 の MAXDOP 設定を変更する必要があります。

何をすべきでしょうか？

- A. db1 に接続し、sp_configure コマンドを実行します。
- B. server1 のマスター データベースに接続し、sp_configure コマンドを実行します。
- C. db1 の拡張プロパティを構成します。
- D. db1 を変更します。

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/configure-max-degree-of-parallelism> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why D is Correct: Modify of db1. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. This fits the case because the requested outcome is: You need to modify the MAXDOP settings for db1. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Connect to db1 and run the sp_configure command. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why B is Wrong: Connect to the master database of server1 and run the sp_configure command. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why C is Wrong: Configure the extended properties of db1. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

質問: 11

MySQLデータベース用のバージョン8.0のAzureデータベースがあります。

どのデータベースクエリが最も多くのリソースを消費するかを特定する必要があります。

どのツールを使用する必要がありますか？

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

正解: [\(正解を表示します\)](#)

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>

質問: 12

SQLMI1 と SQLMI2 という 2 つの Azure SQL マネージド インスタンスを含む Azure サブスクリプションがあります。SQLM2 には、DB1 というデータベースと User1 というユーザーが含まれています。

User1 は DB1 を削除します。

DB1 から SQLMI2 へのポイントインタイム復元を実行する必要があります。

- A. Azure CLI
- B. Transact-SQL
- C. Azureポータル
- D. Azure PowerShell

正解: C ([コメントを发表する](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: The Azure portal is correct because it is the feature whose normal purpose matches the stated requirement. The Azure portal is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to perform a point-in-time restore of DB1 to SQLMI2.

Why A is Wrong: Azure CLI is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: Transact-SQL is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why D is Wrong: Azure PowerShell is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

質問: 13

DB1 という名前の 2 TB Microsoft SQL Server 2019 データベースを含むオンプレミス データセンターがあります。

DB1 を Azure SQL マネージド インスタンスに移行するソリューションを推奨する必要があります。ソリューションでは、ダウンタイムと管理労力を最小限に抑える必要があります。

推奨事項には何を含めるべきですか？

- A. ログ再生サービス (LRS)
- B. ログ配布

C. トランザクション レプリケーション

D. SQL データ同期

正解: ([正解を表示します](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why A is Correct: Log Replay Service (LRS) matches the expected DP-300 administration action. Log Replay Service replays full, differential, and transaction log backups into Azure SQL Managed Instance, making it appropriate for low-downtime migrations from SQL Server. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to recommend a solution to migrate DB1 to an Azure SQL managed instance.

Why B is Wrong: Log shipping is a SQL Server HA/DR technique that copies and restores log backups to a secondary database; it is more manual and operationally heavier for Azure SQL Managed Instance migration than LRS. It can move data in some architectures, but it either increases manual administration or does not match a managed migration into Azure SQL Managed Instance.

Why C is Wrong: Transactional replication distributes committed changes to subscribers, but it is not the default full-database migration method for Azure SQL Managed Instance when the requirement is minimal effort. It can move data in some architectures, but it either increases manual administration or does not match a managed migration into Azure SQL Managed Instance.

Why D is Wrong: SQL Data Sync synchronizes selected data between Azure SQL databases and SQL Server, but it is not a full managed-instance migration mechanism. It can move data in some architectures, but it either increases manual administration or does not match a managed migration into Azure SQL Managed Instance.

Verification Note: The original listed answer was B; the corrected answer is A because Log Replay Service is designed for low-downtime migration to Azure SQL Managed Instance from SQL Server backups with less administrative effort than manual log shipping.

質問: 14

MI1という名前のAzureSQLマネージドインスタンスがあります。

MI1のデータベースの自動チューニングを実装する必要があります。

あなたは何をすべきか？

A. REST APIを使用してパッチ操作を呼び出し、AutomaticTuningServerModeプロパティを変更します。

B. Transact-SQLを使用して、force_last_good_planオプションを有効にします。

C. Azureポータルから、自動チューニングを構成します。

正解: ([正解を表示します](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why B is Correct: Use Transact-SQL to enable the force_last_good_plan option. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to implement automatic tuning for the databases of MI1. Use Transact-SQL to enable the force_last_good_plan option. satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Use the REST API to call the patch operation and modify the AutomaticTuningServerMode property. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Automatic tuning uses built-in recommendations and verification to create/drop indexes or force last known good plans where supported. It does not expose the required metric, query history, wait

/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 15

オンプレミスの Microsoft SQL Server データベースのディザスター リカバリー ソリューションを推奨する必要があります。ソリューションは、次の要件を満たす必要があります。

* 異なる地域へのリアルタイム データ レプリケーションをサポートします。

* Azure をディザスター リカバリーのターゲットとして使用します。

* コストと管理作業を最小限に抑えます。

推奨事項には何を含める必要がありますか？

- A. Azure Virtual Machines 上の SQL Server のインスタンスでのデータベース ミラーリング
- B. Azure Virtual Machines 上の SQL Server の可用性グループ
- C. Azure SQL Managed Instance リンク
- D. Azure SQL Managed Instance へのトランザクション レプリケーション

正解: ([正解を表示します](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: transactional replication to an Azure SQL Managed Instance matches the expected DP-300 administration action. Transactional replication distributes committed changes to subscribers, but it is not the default full-database migration method for Azure SQL Managed Instance when the requirement is minimal effort. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to recommend a disaster recovery solution for an on-premises Microsoft SQL Server database.

Why A is Wrong: database mirroring on an instance of SQL Server on Azure Virtual Machines is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: availability groups for SQL Server on Azure Virtual Machines is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why C is Wrong: an Azure SQL Managed Instance link is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: 16

DB1という名前のデータベースを持つAzure仮想マシン上のSQLServerのインスタンスがあります。

DB1用のAzureSQL DataSyncを実装する予定です。

どの分離レベルを構成する必要がありますか？

- A. シリアル化可能
- B. スナップショット
- C. コミットされていない読み取り
- D. READ COMMITTED

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-data-sync-data-sql-server-sql-database> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: SNAPSHOT matches the expected DP-300 administration action. SNAPSHOT is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Which isolation level should you configure?

Why A is Wrong: SERIALIZABLE is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: READ UNCOMMITTED is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: READ COMMITTED is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: **17**

VNet1という名前の仮想ネットワーク上にVM1という名前のAzure仮想マシンがあります。VM1からインターネットへのアウトバウンドトラフィックはブロックされます。

SqlSrv1という名前の論理サーバー上にSqlDb1という名前のAzureSQLデータベースがあります。

次の要件を満たすには、VM1とSqlDb1の間に接続を実装する必要があります。

VM1がSqlSrv1以外のAzureSQLServerに接続できないことを確認します。

ネットワーク接続をSqlSrv1に制限します。

VNet1で何を作成する必要がありますか？

- A. VPNゲートウェイ
- B. サービスエンドポイント
- C. プライベートリンク
- D. ExpressRouteゲートウェイ

正解: ([正解を表示します](#))

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> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: A service endpoint secures traffic to an Azure service over the Azure backbone but still targets the service public endpoint rather than giving it a private IP in the VNet. This fits the case because the requested outcome is: 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. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: a VPN gateway is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: a private link is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: an ExpressRoute gateway is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

質問: **18**

ファイルがAzureData Lake Storage Gen2コンテナに到着したときに、Azure DataFactoryパイプラインをトリガーする必要があります。

どのリソースプロバイダーを有効にする必要がありますか？

- A. Microsoft.EventHub
- B. Microsoft.EventGrid

C. Microsoft.Sql

D. Microsoft.Automation

正解: ([正解を表示します](#))

Event-driven architecture (EDA) is a common data integration pattern that involves production, detection, consumption, and reaction to events. Data integration scenarios often require Data Factory customers to trigger pipelines based on events happening in storage account, such as the arrival or deletion of a file in Azure Blob Storage account. Data Factory natively integrates with Azure Event Grid, which lets you trigger pipelines on such events.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-event-trigger> Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why B is Correct: Microsoft.EventGrid is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. This fits the case because the requested outcome is: You need to trigger an Azure Data Factory pipeline when a file arrives in an Azure Data Lake Storage Gen2 container. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Microsoft.EventHub is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: Microsoft.Sql is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: Microsoft.Automation is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

質問: 19

SSISDB データベースを含む Azure 仮想マシンに SQL Server 2019 がインストールされている。

最近の障害により、マスター データベースが失われました。

すべての Microsoft SQL Server 統合サービス (SSIS) パッケージが仮想マシン上で実行できないことがわかります。

問題を解決するには、どの 4 つのアクションを順番に実行する必要がありますか? 回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions

Add a certificate to an Azure key vault

Enable Transparent Data Encryption (TDE)

Encrypt a copy of the master key by using the service master key

Turn on the TRUSTWORTHY property and the CLR property

Attach the SSISDB database

Open the master key for the SSISDB database

Answer Area



正解:

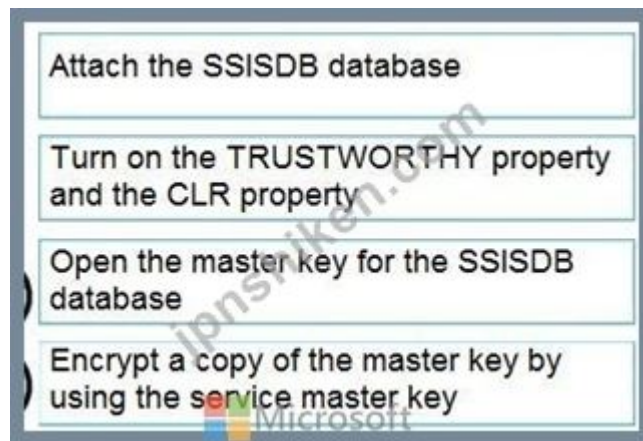
Diagram illustrating the correct sequence of actions for the SSISDB database setup:

Actions	Answer Area
Add a certificate to an Azure key vault	Attach the SSISDB database
Enable Transparent Data Encryption (TDE)	Turn on the TRUSTWORTHY property and the CLR property
Encrypt a copy of the master key by using the service master key	Open the master key for the SSISDB database
Turn on the TRUSTWORTHY property and the CLR property	
Attach the SSISDB database	Encrypt a copy of the master key by using the service master key
Open the master key for the SSISDB database	

The correct sequence of actions is indicated by the dashed red box in the Answer Area, which includes:

- Attach the SSISDB database
- Turn on the TRUSTWORTHY property and the CLR property
- Open the master key for the SSISDB database
- Encrypt a copy of the master key by using the service master key

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>

質問: 20

SOL1 という名前の Azure SQL データベースがあります。

SQL1 の災害復旧ソリューションを実装する必要があります。ソリューションでは、以下の点を最小限に抑える必要があります。

* リカバリポイント目標(JRPO)

* 目標復旧時間(RTO)

* 管理上の努力

ソリューションには何を含めるべきですか？

A. 自動フェイルオーバーグループ

B. Azure サイトの回復

C. 可用性グループ

D. アクティブなジオレプリケーション

正解: A ([コメントを发表する](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why A is Correct: auto failover groups is correct because it is the feature whose normal purpose matches the stated requirement. Auto-failover groups provide managed geo-replication and listener-based failover for Azure SQL Database or Managed Instance, including automatic failover policies for regional outages. The scenario wording points to that specific behavior: You need to implement a disaster recovery solution for SQL1.

Why B is Wrong: Azure Site Recovery is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: availability groups is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why D is Wrong: active geo - replication is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: 21

Azure Virtual Machines インスタンス VM1 上に SQL Server があります。次のクエリを実行します。

```
BACKUP LOG DB1 TO DISK = '\\File1\SQLBackups\DB1.trn'
WITH NORECOVERY,COPY_ONLY,CONTINUE_AFTER_ERROR;
GO
```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

Answer Area

Statements	Yes	No
The log file will be truncated.	☐	☐
DB1 will be placed in an offline state.	☐	☐
You are performing a tail-log backup.	☐	☐

正解:

Statements	Yes	No
The log file will be truncated.	☒	☐
DB1 will be placed in an offline state.	☐	☒
You are performing a tail-log backup.	☐	☒

Explanation:

Answer Area

Statements	Yes	No
The log file will be truncated.	☒	☐
DB1 will be placed in an offline state.	☐	☒
You are performing a tail-log backup.	☐	☒

Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have a SQL Server on Azure Virtual Machines instance named VM1 .

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 22

5 つのデータベースをホストする Microsoft SQL Server 2017 サーバーがあります。

データベースを Azure に移行する予定です。

次の要件を満たすソリューションを推奨する必要があります。

* ワークロードの需要に応じてコンピューティングを自動的にスケーリングします

* 秒単位の課金を提供

推奨事項には何を含めるべきですか? 回答するには、回答エリアで適切なオプションを選択してください。

Azure service:

- A single Azure SQL database in the provisioned compute tier
- A single Azure SQL database in the serverless compute tier
- An Azure SQL Database elastic pool
- Azure SQL Managed Instance

Service tier:

- Basic
- General Purpose
- Standard

Microsoft

正解:

Azure service:


- A single Azure SQL database in the provisioned compute tier
- A single Azure SQL database in the serverless compute tier
- An Azure SQL Database elastic pool
- Azure SQL Managed Instance

Service tier:

- Basic
- General Purpose
- Standard

Microsoft

Explanation:



Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to recommend a solution that meets the following requirements: Automatically scales compute based on the workload demand Provides per-second billing What should you include in the recommendation?

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 23

Azure StreamAnalyticsジョブがあります。

ジョブに十分なストリーミングユニットがプロビジョニングされていることを確認する必要があります。

SU%使用率メトリックの監視を構成します。

監視する必要がある2つの追加メトリックはどれですか？それぞれの正解は、解決策の一部を示しています。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. 遅延入力イベント
- B. 故障したイベント
- C. バックログされた入力イベント
- D. 透かしの遅延
- E. 関数イベント

正解: C,D ([コメントを发表する](#))

To react to increased workloads and increase streaming units, consider setting an alert of 80% on the SU Utilization metric. Also, you can use watermark delay and backlogged events metrics to see if there is an impact.

Note: Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job, by increasing the SUs.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring>


質問: 24

General Purpose サービス レベルに DB 1 という名前の Azure SQL データベースがあります。


SQL Insights を使用して DB 1 を監視する必要があります。

解決策には何を含めるべきですか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:

Box 1 = Azure Monitor Agent

Box 2 = An Azure SQL database

<https://docs.microsoft.com/en-us/azure/azure-sql/database/sql-database-paas-overview?view=azuresql> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to monitor DB 1 by using SQL Insights.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

質問: 25

Azure Data Lake Storage Gen2 コンテナがあります。

データはコンテナに取り込まれ、データ統合アプリケーションによって変換されます。その後、データは変更されません。ユーザーはコンテナ内のファイルを読み取ることはできますが、変更することはできません。

次の要件を満たすデータ アーカイブ ソリューションを設計する必要があります。

* 新しいデータは頻繁にアクセスされるため、できるだけ早く利用できる必要があります。

* 5 年以上前のデータはアクセス頻度は低いですが、要求された場合は 1 秒以内に利用可能である必要があります。

* 7 年以上経過したデータにはアクセスしません。7 年経過後は、可能な限り低コストでデータを永続化する必要があります。

* 必要な可用性を維持しながらコストを最小限に抑える必要があります。

データをどのように管理すればよいですか？ 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Five-year-old data: ▼

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data: ▼

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

正解:

Five-year-old data: ▼

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data: ▼

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Explanation:

Five-year-old data: ▼

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data: ▼

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Box 1: Move to cool storage

The cool access tier has lower storage costs and higher access costs compared to hot storage. This tier is intended for data that will remain in the cool tier for at least 30 days. Example usage scenarios for the cool access tier include:

Short-term backup and disaster recovery

Older data not used frequently but expected to be available immediately when accessed

Large data sets that need to be stored cost effectively, while more data is being gathered for future processing

Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Move to archive storage

Example usage scenarios for the archive access tier include:

Long-term backup, secondary backup, and archival datasets

Original (raw) data that must be preserved, even after it has been processed into final usable form

Compliance and archival data that needs to be stored for a long time and is hardly ever accessed

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

質問: 26

次の表に示すリソースを含む Azure サブスクリプションがあります。

Name	Type
App1	Azure web app
db1	Azure SQL database in the serverless tier

App1 は、一定期間非アクティブになった後に db1 にアクセスしようとする、一時的な接続エラーとタイムアウトが発生します。即時のコストを考慮せずに、できるだけ早く Appl で発生した問題を解決するために db1 を変更する必要があります。何をしますか？

- A. db1 に割り当てられている vCore の数を増やします。
- B. db1 の自動一時停止遅延を減らします。
- C. db1 の自動一時停止遅延を無効にします。
- D. db1 の自動チューニングを有効にします。

正解: D (コメントを发表する)

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: Automatic tuning uses built-in recommendations and verification to create/drop indexes or force last known good plans where supported. This fits the case because the requested outcome is: You need to modify db1 to resolve the issues experienced by Appl as soon as possible, without considering immediate costs What you do? The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why B is Wrong: Decrease the auto-pause delay for db1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: Disable auto-pause delay for db1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: 27

Azure SQLDatabaseのマネージドインスタンスがあります。
インスタンスでパフォーマンスの問題が発生し始めます。

問題の原因となっているクエリを特定し、クエリの実行プランを取得する必要があります。このソリューションでは、管理作業を最小限に抑える必要があります。
何を使うべきですか？

- A. Azureポータル
- B. 拡張イベント
- C. クエリストア
- D. 動的管理ビュー

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store?view=sql-server-ver15> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: dynamic management views matches the expected DP-300 administration action. dynamic management views is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to identify which query is causing the issue and retrieve the execution plan for the query.

Why A is Wrong: the Azure portal is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Extended Events is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 28

次の Azure Resource Manager テンプレートがあります。


```

...
"variable": {
  "serverName": "azsqlserver0001"
},
"resources": [
  {
    "name": "[variables('serverName')]",
    "type": "Microsoft.Sql/servers",
    "apiVersion": "2019-06-01-preview",
    "location": "[parameters('location')]",
    "properties": {
      "administratorLogin": "[parameters('administratorLogin')]",
      "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
      "version": "12.0"
    },
    "resources": [
      {
        "name": "[concat(variables('serverName'), '/', parameters('databaseName'))]",
        "type": "Microsoft.Sql/servers/databases",
        "apiVersion": "2020-08-01-preview",
        "location": "[parameters('location')]",
        "kind": "v12.0",
        "sku": {
          "name": "Standard",
          "tier": "Standard",
          "capacity": 10
        },
        "dependsOn": [
          "[concat('Microsoft.Sql/servers/', variables('serverName'))]"
        ],
        "properties": {
        },
        "resources": [
        ]
      }
    ]
  }
]
}
...

```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。
 注意: 正しい選択ごとに 1 ポイントが加算されます。

Statements	Yes	No
The template deploys a serverless Azure SQL database.	☐	☐
The template deploys a database to an Azure SQL Database managed instance.	☐	☐
The pricing tier of the database deployment is based on DTUs.	☐	☐

正解:

Statements	Yes	No
The template deploys a serverless Azure SQL database.	☐	☒
The template deploys a database to an Azure SQL Database managed instance.	☐	☒
The pricing tier of the database deployment is based on DTUs.	☒	☐

Explanation:

Statements	Yes	No
The template deploys a serverless Azure SQL database.	☐	☒
The template deploys a database to an Azure SQL Database managed instance.	☐	☒
The pricing tier of the database deployment is based on DTUs.	☒	☐

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/purchasing-models>

<https://docs.microsoft.com/en-us/azure/azure-sql/database/single-database-create-arm-template-quickstart> Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have the following Azure Resource Manager template.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 29

オンプレミス ネットワークには、db1 という名前のデータベースをホストする Microsoft SQL Server 2016 サーバーが含まれています。

Azure サブスクリプションがあります。

db1 を Azure SQL マネージド インスタンスに移行する予定です。

SQL マネージド インスタンスを作成する必要があります。このソリューションでは、インスタンスのディスク レイテンシを最小限に抑える必要があります。

どのサービス レベルを使用する必要がありますか？

- A. ハイパースケール
- B. 汎用
- C. プレミアム
- D. ビジネスクリティカル

正解: ([正解を表示します](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why D is Correct: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. In this scenario, the important constraint is: You need to create the SQL managed instance. Business Critical satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Hyperscale is designed for very large databases and fast scale-out storage architecture, not merely for intermittent CPU bursts. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why B is Wrong: General Purpose is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: Premium is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

質問: 30

Azure SQL データベースに Customer というテーブルが含まれています。Customer には、次の表に示す列があります。

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

Customer_Phone列に動的データマスクを実装する予定です。マスクは以下の要件を満たす必要があります。

- * 各顧客の電話番号の最初の 6 桁をマスクする必要があります。
- * 各顧客の電話番号の下 4 桁が見える必要があります。
- * ハイフンは保持され、表示される必要があります。

動的データ マスクをどのように構成すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

Exposed Prefix:

▼

0

1

3

5

Padding String:

▼

x

xxxxxx

xxx-xxx

xxx-xxx-

x[3]-x[3]

Exposed Suffix:

▼

0

1

3

5

正解:

Exposed Prefix:

	▼
0	
1	
3	
5	

Padding String:

	▼
X	
XXXXXX	
XXX-XXX	
XXX-XXX-	
x[3]-x[3]	

Exposed Suffix:

	▼
0	
1	
3	
5	



Explanation:

Exposed Prefix:

	▼
0	
1	
3	
5	

Padding String:

	▼
X	
XXXXXX	
XXX-XXX	
XXX-XXX-	
x[3]-x[3]	

Exposed Suffix:

	▼
0	
1	
3	
5	



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> Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: The mask must meet the following requirements:

The first six numerals of each customer's phone number must be masked.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 31

AzSQL1 という名前の Azure SQL サーバーに、DB1 という名前の新しい Azure SQL データベースがあります。

作成された唯一のユーザーはサーバー管理者です。

認証に Azure Active Directory (Azure AD) を使用する包含データベース ユーザーを DB1 に作成する必要があります。

順番に実行する必要がある 3 つのアクションはどれですか。回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

The screenshot shows a task interface with two main sections: "Actions" and "Answer Area".

Actions:

- Connect to DB1 by using the Active Directory admin account.
- Create a user by using the FROM EXTERNAL PROVIDER clause.
- Connect to DB1 by using the server administrator account.
- Set the Active Directory Admin for AzSQL1.
- From the Azure portal, assign the SQL DB Contributor role to the user.
- Create a login in the master database.

Answer Area:

The Answer Area contains a Microsoft logo and four circular buttons with arrows: a right arrow, a left arrow, an up arrow, and a down arrow. A large, faint watermark "jnshtken.com" is visible across the center of the interface.

正解:

Actions

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Connect to DB1 by using the server administrator account.

Set the Active Directory Admin for AzSQL1.

From the Azure portal, assign the SQL DB Contributor role to the user.

Create a login in the master database.

Answer Area

Set the Active Directory Admin for AzSQL1.

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

Explanation:

Set the Active Directory Admin for AzSQL1.

Connect to DB1 by using the Active Directory admin account.

Create a user by using the FROM EXTERNAL PROVIDER clause.

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>

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: **32**

次のデータベースレベルおよびインスタンスレベルの機能を使用するDB1という名前のMicrosoftSQL Server2019データベースがあります。

クラスター化された列ストアインデックス

自動チューニング

変更の追跡

PolyBase

DB1をAzureSQLデータベースに移行することを計画しています。

DB1を移行する前に、どの機能を削除または置換する必要がありますか？

A. クラスター化された列ストアインデックス

B. PolyBase

C. 変更追跡

D. 自動チューニング

正解: ([正解を表示します](#))

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

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/polybase/polybase-versioned-feature-summary> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: PolyBase is correct because it is the feature whose normal purpose matches the stated requirement. PolyBase is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: You have a Microsoft SQL Server 2019 database named DB1 that uses the following database-level and instance-level features.

Why A is Wrong: Clustered columnstore indexes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Change tracking is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: Automatic tuning uses built-in recommendations and verification to create/drop indexes or force last known good plans where supported. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 33

AzureSQLデータベースを使用するソリューションを計画しています。ソリューションの使用は、毎年10月1日から1月1日までピークになります。

ピーク使用時には、データベースには次のものが必要になります。

24コア

500GBのストレージ

124GBのメモリ

50,000IOPS以上

オフピークの使用期間中、AzureSQLデータベースのサービス層は標準に設定されます。

ピーク使用時にどのサービス層を使用する必要がありますか？

A. ビジネスクリティカル

B. プレミアム

C. ハイパースケール

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-vcost-single-databases#business-critical---provisioned-compute---gen4> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why A is Correct: Business Critical matches the expected DP-300 administration action. Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You are planning a solution that will use Azure SQL Database.

Why B is Wrong: Premium is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Hyperscale is designed for very large databases and fast scale-out storage architecture, not merely for intermittent CPU bursts. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

質問: **34**

タスク12

db1には高可用性を構成する必要があります。ソリューションは、Azure データセンターの障害を許容し、データの損失やアプリケーションの接続文字列の変更を必要としない必要があります。

正解:

See the explanation part for the complete Solution.

Explanation:

To configure high availability for db1, you can use the failover groups feature of Azure SQL Database. Failover groups allow you to manage the replication and failover of a group of databases across different regions with the same connection strings¹. You can choose all, or a subset of, user databases in a logical server to be replicated to another logical server in a different region. You can also specify the failover policy, such as manual or automatic, and the grace period for data loss.

Here are the steps to create a failover group for db1:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

Select Failover groups in the left menu and click on Add group.

Enter a name for the failover group and select a secondary region that is different from the primary region.

Click on Create a new server and enter the details for the secondary server, such as server name, admin login, password, and subscription.

Click on Select existing database(s) and choose db1 from the list of databases on the primary server.

Click on Configure failover policy and select the failover mode, grace period, and read-write failover endpoint mode according to your preferences.

Click on Create to create the failover group and start the replication of db1 to the secondary server.

Using PowerShell commands:

Install the Azure PowerShell module and log in with your Azure account.

Run the following command to create a new server in the secondary region: `New-AzSqlServer - ResourceGroupName < your-resource-group-name > -ServerName < your-secondary-server-name > - Location " < secondary-region-name > " -SqlAdministratorCredentials $(New-Object -TypeName System.`

`Management.Automation.PSCredential -ArgumentList " < your-admin-login > " , $(ConvertTo-SecureString - String " < your-password > " -AsPlainText -Force))` Run the following command to create a new failover group with db1: `New-AzSqlDatabaseFailoverGroup - ResourceGroupName < your-resource-group-name > -ServerName < your-primary-server-name > - PartnerResourceGroupName < your-resource-group-name > -PartnerServerName < your-secondary-server-name > -FailoverGroupName < your-failover-group-name > -Database db1 -FailoverPolicy < manual-or- automatic > - GracePeriodWithDataLossHours < grace-period-in-hours > -ReadWriteFailoverEndpoint " < enabled-or-disabled > "` You can modify the parameters of the command according to your preferences, such as the failover policy, grace period, and read-write failover endpoint mode.

These are the steps to create a failover group for db1

質問: 35

db1 という名前のデータベースがあります。

db1 のログには次のエントリが含まれています。

Date 10/5/2021 10:57:08 AM

Log SQL Server (Current - 10/5/2021 11:26:00 AM)

Source spid1595

Message

The transaction log for database 'db1' is full due to 'AVAILABILITY_REPLICA'

db1 がトランザクションを処理できることを確認する必要があります。

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

Answer Area

正解:

Actions

- Add db1 back to the availability group.
- Shrink db1.
- Shrink the transaction log file.
- Remove db1 from the availability group.
- Back up the transaction log file.

Answer Area

- Remove db1 from the availability group.
- Shrink the transaction log file.
- Add db1 back to the availability group.

Explanation:

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure That db1 can process transactions.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 36

DB1という名前のデータベースをホストするオンプレミスのMicrosoftSQL Server2019サーバーがあります。

SQLMI1という名前のAzureSQLマネージドインスタンスとVNET1という名前の仮想ネットワークを含むAzureサブスクリプションがあります。SQLMI1はVNET1に常駐します。オンプレミスネットワークは、ExpressRoute接続を使用してVNET1に接続します。

Azure Database Migration Serviceを使用して、DB1をSQLMI1に移行することを計画しています。

移行をサポートするようにVNET1を構成する必要があります。

あなたは何をするべきか？

- A. サービスエンドポイントを構成します。
- B. 仮想ネットワークピアリングを構成します。
- C. Azureファイアウォールをデプロイします。
- D. ネットワークセキュリティグループ (NSG)を構成します。

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-to-managed-instance> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why A is Correct: A service endpoint secures traffic to an Azure service over the Azure backbone but still targets the service public endpoint rather than giving it a private IP in the VNet. The scenario asks for: You need to configure VNET1 to support the migration. That makes Configure service endpoints. the option that satisfies the required Azure SQL layer and operational outcome.

Why B is Wrong: Configure virtual network peering. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Deploy an Azure firewall. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: Configure network security groups (NSGs). is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

質問: 37

販売データのストリーミング集計を実行するには、どのウィンドウ関数を使用する必要がありますか？


- A. スライディング
- B. ホッピング
- C. セッション
- D. タンブリング

正解: ([正解を表示します](#))

Scenario: The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping.

Tumbling window functions are used to segment a data stream into distinct time segments and perform a function against them, such as the example below. The key differentiators of a Tumbling window are that they repeat, do not overlap, and an event cannot belong to more than one tumbling window.

Tell me the count of Tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second, 10)
```

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/stream-analytics/stream-analytics-window-functions.md>

質問: 38

Azure サブスクリプションをお持ちです。

書き込みアクセラレータをサポートする Azure Virtual Machines に SQL Server のインスタンスを展開する予定です。

どの仮想マシン シリーズを使用すればよいですか？

- A. Hシリーズ
- B. Gシリーズ
- C. Mシリーズ

D. Eシリーズ

正解: ([正解を表示します](#))

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: M-series matches the expected DP-300 administration action. M-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Which virtual machine series should you use?

Why A is Wrong: H-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why B is Wrong: G-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.


Why D is Wrong: E-series can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

質問: 39

Table1 という名前のテーブルを含む、という名前の Azure SQL データベースがあります。

Table1 に不正なデータを入力するクエリを実行します。

ロード操作中の表 1 のパフォーマンスを図に示します。



正解:

To reduce how long it takes to complete the query you must [answer choice].

scale the resource
use an elastic pool
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

a temporary I
an In-Memory OTLP durable
an In-Memory OTLP non durable

Microsoft

Explanation:

To reduce how long it takes to complete the query you must [answer choice].

scale the resource
use an elastic pool
perform query tuning

To reduce the log IO load of the operation, the query must be updated to use [answer choice] table.

a temporary I
an In-Memory OTLP durable
an In-Memory OTLP non durable

Microsoft

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure SQL database named that contains a table named Table1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 40

10 の Azure リージョンにまたがる 100 の Azure SQL マネージド インスタンスを管理します。

メンテナンス イベントが 10 のリージョンのいずれかに影響を与える場合、ボイス メッセージ通知を受信する必要があります。ソリューションは、管理作業を最小限に抑える必要があります。

あなたは何をすべきか？

- A. Azure portal から、サービス正常性アラートを作成します。
- B. Azure portal から、Azure Advisor オペレーショナル エクセレンス アラートを作成します。
- C. Microsoft SQL Server Management Studio (SSMS) から、SQL Server エージェント ジョブを構成します。

D. Azure portal から、アクティビティ ログ アラートを構成します。

正解: ([正解を表示します](#))

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why A is Correct: From the Azure portal, create a service health alert. matches the expected DP-300 administration action. From the Azure portal, create a service health alert. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to receive voice message notifications when a maintenance event affects any of the 10 regions.

Why B is Wrong: Microsoft Entra ID authentication allows database users to authenticate with organizational identities instead of SQL logins, supporting centralized identity governance. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: SQL Server Agent schedules and runs SQL Server jobs; in Azure SQL Database, Elastic Jobs or Azure Automation are often used instead. It is not wrong technology in general, but it is the wrong HA

/DR control for this scenario ' s failure model.

Why D is Wrong: From the Azure portal, configure an activity log alert. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: **41**

2つの100GBデータベースをAzureに移動することを計画しています。

ワークロードに基づいてリソース消費を動的にスケーリングする必要があります。ソリューションは、スケーリング操作中のダウンタイムを最小限に抑える必要があります。

何を使うべきですか？

A. エラスティックプール内の2つのAzureSQLデータベース

B. Azure仮想マシン上のSQLServerでホストされている2つのデータベース

C. AzureSQLマネージドインスタンス内の2つのデータベース

D. 2つの単一のAzureSQLデータベース

正解: ([正解を表示します](#))

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> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why A is Correct: two Azure SQL Databases in an elastic pool matches the expected DP-300 administration action. Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to dynamically scale resources consumption based on workloads.

Why B is Wrong: two databases hosted in SQL Server on an Azure virtual machine is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: two databases in an Azure SQL Managed instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: two single Azure SQL databases is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

質問: **42**

タスク1

db1 という名前の Azure SQL データベースで、SalesLT.SalesOrderHeader テーブルの PK_SalesOrderHeader_SalesOrderID クラスター化インデックスでページ圧縮を有効にする必要があります。

正解:

See the explanation part for the complete Solution.

Explanation:

To enable page compression on the PK_SalesOrderHeader_SalesOrderID clustered index of the SalesLT.

SalesOrderHeader table in db1, you can use the following Transact-SQL script:

```
-- Connect to the Azure SQL database named db1
```

```
USE db1;
```

```
GO
```

```
-- Enable page compression on the clustered index
```

```
ALTER INDEX PK_SalesOrderHeader_SalesOrderID ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO
```

This script will rebuild the clustered index with page compression, which can reduce the storage space and improve the query performance. The script solution consists of three parts:

The first part is USE db1; GO. This part connects to the Azure SQL database named db1, where the SalesLT.

SalesOrderHeader table is located. The GO command separates the batches of Transact-SQL statements and sends them to the server.

The second part is ALTER INDEX PK_SalesOrderHeader_SalesOrderID ON SalesLT.SalesOrderHeader REBUILD WITH (DATA_COMPRESSION = PAGE); GO. This part enables page compression on the clustered index named PK_SalesOrderHeader_SalesOrderID, which is defined on the SalesLT.

SalesOrderHeader table. The ALTER INDEX statement modifies the properties of an existing index.

The REBUILD option rebuilds the index from scratch, which is required to change the compression setting.

The DATA_COMPRESSION = PAGE option specifies that page compression is applied to the index, which means that both row and prefix compression are used. Page compression can reduce the storage space and improve the query performance by compressing the data at the page level. The GO command ends the batch of statements.

The third part is optional, but it can be useful to verify the compression status of the index. It is SELECT name, index_id, data_compression_desc FROM sys.indexes WHERE object_id = OBJECT_ID(' SalesLT.

SalesOrderHeader ');. This part queries the sys.indexes catalog view, which contains information about the indexes in the database. The SELECT statement returns the name, index_id, and data_compression_desc columns for the indexes that belong to the SalesLT.SalesOrderHeader table. The OBJECT_ID function returns the object identification number for the table name. The data_compression_desc column shows the compression type of the index, which should be PAGE for the clustered index after the script is executed.

These are the steps of the script solution for enabling page compression on the clustered index of the SalesLT.

SalesOrderHeader table in db1.

質問: 43

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、述べられた目標を達成する可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

Salesという名前のAzureSQLデータベースがあります。

次の要件を満たすには、Salesのディザスタリカバリを実装する必要があります。

*通常の操作中に、Salesの読み取り可能なコピーを少なくとも2つ提供します。

*データセンターに障害が発生した場合でも、セールスを引き続き利用できるようにします。

解決策：汎用サービス層とジオレプリケーションを使用するAzureSQLデータベースをデプロイします。

これは目標を達成していますか？

A. はい

B. いいえ

正解: [\(正解を表示します\)](#)

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>

質問: 44

次の表に示すリソースを含む Azure サブスクリプションがあります。

Name	Type	Azure region
VM1	Azure virtual machine	West US 2
MI1	Azure SQL Managed Instance	East US

VM1 と MIL 間の接続を構成する必要があります。ソリューションは次の要件を満たしている必要があります。

* 接続は暗号化されている必要があります。

* ネットワーク遅延を最小限に抑える必要があります。

何を実装すべきでしょうか？

- A. 仮想ネットワークピアリング
- B. プライベートエンドポイント
- C. サービスエンドポイント
- D. サイト間VPN

正解: [\(正解を表示します\)](#)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B is Correct: A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. The scenario asks for: You need to configure a connection between VM1 and MIL The solution must meet the following requirements: * The connection must be encrypted. That makes private endpoints the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: virtual network peering is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to configure a connection between VM1 and MIL The solution must meet the following requirements: * The connection must be encrypted.

Why C is Wrong: A service endpoint secures traffic to an Azure service over the Azure backbone but still targets the service public endpoint rather than giving it a private IP in the VNet. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: a site-to-site VPN is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to configure a connection between VM1 and MIL The solution must meet the following requirements: * The connection must be encrypted.

質問: 45

contoso.comという名前のドメインを使用するAzureサブスクリプションがあります。

DBServer1とDBServer2という名前の2つのAzureVMがあります。それぞれがデフォルトのSQLServerインスタンスをホストします。DBServer1は米国東部のAzureリージョンにあり、DatabaseAという名前のデータベースが含まれています。

DBServer2は、米国西部のAzureリージョンにあります。

DBServer1には、大量のデータ変更とデータ書き込みの低遅延要件があります。

DatabaseAの新しい可用性グループを構成する必要があります。セカンダリレプリカはDBServer2に常駐します。

あなたは何をするべきか？

- A. プライマリエンドポイントをTCP ://DBServer1.contoso.com :445として構成し、セカンダリエンドポイントをTCP ://DBServer2.contoso.com :445として構成し、可用性モードを非同期に設定します。
- B. プライマリエンドポイントをTCP ://DBServer1.contoso.com :445として構成し、セカンダリエンドポイントをTCP ://DBServer2.contoso.com :445として構成し、可用性モードを同期に設定します。
- C. プライマリエンドポイントをTCP ://DBServer1.contoso.com :5022として構成し、セカンダリエンドポイントをTCP ://DBServer2.contoso.com :5022として構成し、可用性モードを非同期に設定します。
- D. プライマリエンドポイントをTCP ://DBServer1.contoso.com :5022として構成し、セカンダリエンドポイントをTCP ://DBServer2.contoso.com :5022として構成し、可用性モードを同期に設定します。

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/availability-modes-always-on-availability-groups?view=sql-server-ver15> Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why C is Correct: Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Asynchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. This fits the case because the requested outcome is: DBServer1 has a high volume of data changes and low latency requirements for data writes. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Asynchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why B is Wrong: Configure the primary endpoint as TCP://DBServer1.contoso.com:445, configure the secondary endpoint as TCP://DBServer2.contoso.com:445, and set the availability mode to Synchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: Configure the primary endpoint as TCP://DBServer1.contoso.com:5022, configure the secondary endpoint as TCP://DBServer2.contoso.com:5022, and set the availability mode to Synchronous. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

質問: 46

Azure SQL マネージド インスタンスがあります。

Transact-SQL を使用して DB1 という名前のデータベースを復元する必要があります。

どのコマンドを実行する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

RESTORE ▼ DB1 FROM ▼

DATABASE
FILE
LOG

DISK = N'\\NAS01\SQLBackups\DB1.bak';
TAPE = N'\\Tape0'
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

正解:

RESTORE ▼ DB1 FROM ▼

DATABASE
FILE
LOG

DISK = N'\\NAS01\SQLBackups\DB1.bak';
TAPE = N'\\Tape0'
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Explanation:

RESTORE ▼ DB1 FROM ▼

DATABASE
FILE
LOG

DISK = N'\\NAS01\SQLBackups\DB1.bak';
TAPE = N'\\Tape0'
URL = N'https://mybackups.blob.core.windows.net/bkups/DB1.bak'

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to restore a database named DB1 by using Transact-SQL.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> 503問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 47

タスク4

db1 の SalesLT.Customer テーブルの LastName 列を、Always Encrypted を使用して暗号化する必要があります。

Windows証明書ストアを使用する必要があります。SQL Server Management StudioとAzureポータルを使用する必要がある場合もあります。

正解:

See the explanation part for the complete Solution.

Explanation:

Use SQL Server Management Studio Always Encrypted Wizard and encrypt:

Database: db1

Schema: SalesLT

Table: Customer

Column: LastName

Encryption feature: Always Encrypted

Column master key store: Windows Certificate Store

Do not use Transparent Data Encryption. TDE encrypts the database at rest, but this task specifically requires Always Encrypted, which protects selected columns and keeps the encryption keys outside the database engine. Microsoft states that Always Encrypted uses column encryption keys to encrypt column data and column master keys to protect those column encryption keys. Column master keys are stored outside the database system, such as in the Windows certificate store.

Method 1 - SSMS Always Encrypted Wizard

This is the correct method for the simulation.

Step 1: Open SSMS and connect to Azure SQL Database

Open SQL Server Management Studio.

Connect to the Azure SQL logical server that hosts db1.

Use a SQL admin account or Microsoft Entra admin account.

In Options > Connection Properties, select database:

db1

Connect.

If SSMS cannot connect, go to the Azure portal and add your client IP address under the SQL server firewall

/networking settings.

Step 2: Open the table column

In Object Explorer:

Expand Databases.

Expand db1.

Expand Tables.

Expand:

SalesLT.Customer

Expand Columns.

Locate:

LastName

Microsoft confirms that the Always Encrypted Wizard can be launched at the database, table, or individual column level. For one column, the cleanest path is to launch it directly from the column.

Step 3: Launch the Always Encrypted Wizard

Right-click the LastName column, then select:

Encrypt Column

or, depending on the SSMS version:

Always Encrypted Wizard

Alternative path:

Right-click db1 > Tasks > Always Encrypted Wizard

Then manually select:

SalesLT.Customer.LastName

Step 4: Select the LastName column for encryption

On the Column Selection page:

Find:

SalesLT.Customer.LastName

Select the checkbox for LastName.

Set the encryption type.

Use:

Randomized

unless the lab specifically requires searching or equality filtering on LastName.

Reason: Randomized encryption is stronger because identical plaintext values produce different ciphertext values. Deterministic encryption allows equality lookups, joins, grouping, and indexing, but leaks more pattern information because identical plaintext values produce identical encrypted values. Microsoft describes deterministic encryption as query-friendly but more pattern-revealing, while randomized encryption is more secure but does not support normal searching/grouping/joining without secure enclaves.

For this task, the requirement is only to encrypt the LastName column, so Randomized is the safer default.

Step 5: Choose or create a Column Encryption Key

For Encryption Key, select a new key such as:

CEK_Auto1

or create a new column encryption key if one does not already exist.

This is the key that encrypts the data in the LastName column. Microsoft states that a column encryption key encrypts the data in encrypted columns, and the column master key encrypts/protects the column encryption key.

Select Next.

Step 6: Configure the Column Master Key in Windows Certificate Store

On the key configuration page, create a new Column Master Key.

Use settings like:

Setting

Value

Column master key name

CMK_WindowsCert or default generated name

Key store

Windows Certificate Store

Certificate location

Current User

Certificate

Generate new certificate

Column encryption key

CEK_Auto1 or default generated CEK

In many SSMS versions, the wizard creates both:

CMK_Auto1

CEK_Auto1

That is acceptable as long as the CMK key store is Windows Certificate Store.

Microsoft states that SQL Server Management Studio supports column master keys stored in the Windows Certificate Store, and that a column master key can be a certificate stored in Windows Certificate Store.

Step 7: Run the wizard

On the final wizard pages:

Review the configuration.

Choose:

Proceed to finish now

or:

Run immediately

Select Finish.

SSMS will generate the column master key metadata, column encryption key metadata, and perform the data encryption operation. Microsoft explains that the wizard can encrypt selected columns and can generate a new column master key and column encryption key when needed.

During encryption, SSMS may temporarily create a new table, copy data, encrypt the selected column, and swap the table back, depending on whether secure enclaves are used. Microsoft notes that the wizard can move data out of the database and perform cryptographic operations inside the SSMS process when secure enclave in-place encryption is not used.

Verification Steps

Step 1: Confirm Always Encrypted metadata exists

Run this in db1:

```
SELECT  
name,  
key_store_provider_name,  
key_path  
FROM sys.column_master_keys;
```

You should see a column master key that uses the Windows certificate store provider.

Then run:

```
SELECT  
name  
FROM sys.column_encryption_keys;
```

You should see the column encryption key created by the wizard.

Step 2: Confirm LastName is encrypted

Run:

```
SELECT  
t.name AS table_name,  
c.name AS column_name,  
c.encryption_type_desc,  
cek.name AS column_encryption_key  
FROM sys.columns AS c  
JOIN sys.tables AS t  
ON c.object_id = t.object_id  
LEFT JOIN sys.column_encryption_keys AS cek  
ON c.column_encryption_key_id = cek.column_encryption_key_id  
WHERE t.name = ' Customer '  
AND SCHEMA_NAME(t.schema_id) = ' SalesLT '  
AND c.name = ' LastName ' ;
```

Expected result:

table_name: Customer

column_name: LastName

encryption_type_desc: RANDOMIZED

column_encryption_key: CEK_Auto1 or similar

If you selected deterministic encryption, the expected value will be:

DETERMINISTIC

The key requirement is that encryption_type_desc is no longer NULL.

Step 3: Test viewing the encrypted column

Open a new SSMS query connection without Always Encrypted enabled and run:

```
SELECT TOP 10 LastName
```

```
FROM SalesLT.Customer;
```

You should not see normal plaintext values.

Then reconnect with Always Encrypted enabled:

In SSMS, select Connect > Database Engine.

Select Options.

Go to Additional Connection Parameters.

Add:

Column Encryption Setting=Enabled

Connect again.

Run:

```
SELECT TOP 10 LastName
```

```
FROM SalesLT.Customer;
```

A client that has access to the Windows certificate/private key should be able to decrypt the values. Microsoft explains that the database stores encrypted data and key metadata, while client-side components with access to the column master key perform encryption and decryption.

Important Exam Notes

Do not choose Azure Key Vault

The task explicitly says:

You must use the Windows Certificate Store.

So the column master key should not be stored in Azure Key Vault.

Wrong:

Azure Key Vault

Correct:

Windows Certificate Store

Do not use TDE

TDE is not column-level Always Encrypted. It encrypts database files/logs at rest, but users and administrators querying the database still see plaintext if they have SQL permissions.

Correct technology:

Always Encrypted

Correct key store:

Windows Certificate Store

Correct target column:

SalesLT.Customer.LastName

質問: 48

5つのデータベースをホストするSQL1という名前のオンプレミスのMicrosoftSQLServerがあります。

データベースをAzureSQLマネージドインスタンスに移行する必要があります。このソリューションでは、ダウンタイムを最小限に抑え、データの損失を防ぐ必要があります。

何を使うべきですか？

A. ログ配布

B. 常時オンの可用性グループ

C. データベース移行アシスタント

D. バックアップと復元

正解: C ([コメントを发表する](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why C is Correct: Database Migration Assistant is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario asks for: You need to migrate the databases to an Azure SQL managed instance. That makes Database Migration Assistant the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Log shipping is a SQL Server HA/DR technique that copies and restores log backups to a secondary database; it is more manual and operationally heavier for Azure SQL Managed Instance migration than LRS. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why B is Wrong: Always On availability groups are SQL Server/SQL Managed Instance HA technology, not the Azure SQL Database PaaS failover abstraction used for database groups. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: Backup and Restore is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

質問: 49

タスク1

db1 という名前の Azure SQL データベースに対して、アクティブ ジオ レプリケーションを使用して災害復旧ソリューションを実装する必要があります。レプリカは、sql60152867-

dr.database.windows.net という名前のサーバー上の、米国東部または米国東部 2 Azure リージョンに配置する必要があります。SQL Server Management Studio と Azure ポータルを使用する必要があります。

正解:

See the explanation part for the complete Solution.

Explanation:

Requirement: Configure active geo-replication for Azure SQL Database db1. The geo-replica must be created in East US or East US 2 on the logical Azure SQL server:

sql60152867-dr.database.windows.net

The important point: in Azure SQL, the logical server name used in portal/T-SQL is usually:

sql60152867-dr

not the full FQDN.

Microsoft states that active geo-replication is configured per database, and a geo-secondary is created for an existing Azure SQL Database. After creation and seeding, changes from the primary are replicated asynchronously to the secondary.

Method 1 - Azure Portal Method

This is the safest method for the simulation if the portal is available.

Step 1: Open the primary database

Sign in to the Azure portal.

Search for SQL databases.

Select the database named db1.

Confirm you are looking at the primary database, not an existing secondary.

Step 2: Open the Replicas blade

In the left menu of the db1 database page, scroll to Data management.

Select Replicas.

Select Create replica.

Microsoft's portal workflow is: open the database, go to Data management > Replicas, and choose Create replica.

Step 3: Configure the geo-secondary replica

On the Create SQL Database replica page, configure the target like this:

Setting

Value

Database

db1

Replica type

Geo replica / Active geo-replication

Target server

sql60152867-dr

Region

East US or East US 2

Database name

db1

Compute + storage

Keep same or compatible with primary

Elastic pool

Only choose this if the lab specifically requires an elastic pool

Do not create a failover group unless the task asks for one. This task says active geo replication, so configure a database-level geo-replica, not a failover group. Microsoft explicitly separates active geo-replication from failover groups and notes that active geo-replication is configured per database.

Step 4: Review and create

Select Review + create.

Confirm the target server is:

sql60152867-dr

Confirm the region is either:

East US

or

East US 2

Select Create.

Azure will create the secondary database and begin the seeding process. Microsoft notes that the secondary database has the same name as the primary by default and begins replication after it is created and seeded.

Step 5: Verify replication

After deployment completes:

Go back to the primary database db1.

Open Replicas again.

Under Geo replicas, confirm that a replica exists on:

sql60152867-dr.database.windows.net

Confirm the replica status is healthy, online, or synchronizing.

You can also open the target SQL server sql60152867-dr and verify that a database named db1 now exists there.

Method 2 - SSMS / T-SQL Method

Use this method if the portal is awkward or the exam simulation expects T-SQL.

Step 1: Allow SSMS connectivity

Before connecting with SSMS:

In Azure portal, open the primary SQL server hosting db1.

Go to Networking or Firewalls and virtual networks.

Add your client IP address.

Repeat this on the secondary server:

sql60152867-dr.database.windows.net

This matters because SSMS must be able to connect to the Azure SQL logical server.

Step 2: Connect to the primary server in SSMS

Open SQL Server Management Studio.

Connect to the primary Azure SQL logical server that hosts db1.

Use SQL admin credentials or Microsoft Entra admin credentials.

In Connection Properties, connect to the database:

master

This is important. For Azure SQL Database geo-replication setup through T-SQL, run the command from the master database on the primary server.

Step 3: Run the active geo-replication command

Run this query:

```
ALTER DATABASE [db1]
```

```
ADD SECONDARY ON SERVER [sql60152867-dr]
```

```
WITH (ALLOW_CONNECTIONS = ALL);
```

Microsoft documents that ALTER DATABASE ... ADD SECONDARY ON SERVER creates a secondary database for an existing Azure SQL Database and starts data replication. The official example also uses WITH (ALLOW_CONNECTIONS = ALL) to create a readable geo-secondary.

Step 4: Verify the replication link

Still connected to the primary server, run:

```
SELECT *
```

```
FROM sys.geo_replication_links;
```

Or use:

```
SELECT *
```

```
FROM sys.dm_geo_replication_link_status;
```

Microsoft lists sys.geo_replication_links and sys.dm_geo_replication_link_status as views used to return information about existing replication links and replication status.

Optional PowerShell Method

Only use this if the lab gives you Cloud Shell or PowerShell access.

```
New-AzSqlDatabaseSecondary `
-ResourceGroupName " < PrimaryResourceGroupName > " `
-ServerName " < PrimaryServerName > " `
-DatabaseName " db1 " `
-PartnerResourceGroupName " < SecondaryResourceGroupName > " `
-PartnerServerName " sql60152867-dr " `
-PartnerDatabaseName " db1 "
```

Microsoft identifies New-AzSqlDatabaseSecondary as the PowerShell cmdlet that creates a secondary database for an existing Azure SQL Database and starts replication.

Final Answer / What You Must Achieve

The task is complete when:

Database db1 still exists on the primary server.

A secondary database named db1 exists on:

sql60152867-dr.database.windows.net

The target server is in East US or East US 2.

The replica appears under db1 > Replicas > Geo replicas.

Replication status is healthy, online, seeding, or synchronizing.

You did not configure a failover group unless separately requested.

質問: 50

オンプレミスの Microsoft SQL Server 2022 インスタンス (名前: SQL1) があります。Azure サブスクリプションには、Azure SQL マネージドインスタンス (名前: SQLM11) が含まれています。SQL1 と SQLM11 間で双方向の災害復旧ソリューションを構成する必要があります。ソリューションには何を含めるべきですか？

- A. ログ再生サービス (LRS)
- B. マネージドインスタンスリンク
- C. Azure Database Migration Service
- D. フェイルオーバーグループ

正解: ([正解を表示します](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: the Managed Instance link is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement.

In this scenario, the important constraint is: You have an Azure subscription that contains an Azure SQL managed instance named SQLM11. You need to configure a bidirectional disaster recovery solution between SQL1 and SQLM11. the Managed Instance link satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Log Replay Service replays full, differential, and transaction log backups into Azure SQL Managed Instance, making it appropriate for low-downtime migrations from SQL Server. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: Azure Database Migration Service is the managed migration service used to move database engines to Azure targets with online or offline migration modes depending on source and target support. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: Auto-failover groups provide managed geo-replication and listener-based failover for Azure SQL Database or Managed Instance, including automatic failover policies for regional outages. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

質問: 51

SQL1 という名前の新しい Azure SQL マネージド インスタンスを作成し、データベース メールの拡張ストア プロシージャを有効にします。

SQL 1 で実行されている SOL サーバー エージェント ジョブが、障害が発生したときに管理者に通知できるようにする必要があります。

順番に実行する必要がある 3 つのアクションはどれですか。回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

The screenshot shows the 'Actions' list on the left and the 'Answer Area' on the right. The 'Actions' list contains five items: 'Enable pager notifications upon failure.', 'Create a profile named application_dbmail_profile.', 'Create a Database Mail account.', 'Create a profile named AzureManagedInstance_dbmail_profile.', and 'Enable email notifications upon failure.'. The 'Answer Area' is empty, with a numbered list (1, 2, 3) and arrows indicating the sequence of actions to be added.

正解:

The screenshot shows the 'Actions' list on the left and the 'Answer Area' on the right. The 'Answer Area' contains three actions in the correct sequence: 1. 'Create a Database Mail account.', 2. 'Create a profile named AzureManagedInstance_dbmail_profile.', and 3. 'Enable email notifications upon failure.'. The 'Actions' list on the left shows the remaining actions: 'Enable pager notifications upon failure.', 'Create a profile named application_dbmail_profile.', and 'Create a Database Mail account.'.

Explanation:

The screenshot shows the 'Actions' list on the left and the 'Answer Area' on the right. The 'Answer Area' contains three actions in the correct sequence: 1. 'Create a Database Mail account.', 2. 'Create a profile named AzureManagedInstance_dbmail_profile.', and 3. 'Enable email notifications upon failure.'. The 'Actions' list on the left shows the remaining actions: 'Enable pager notifications upon failure.' and 'Create a profile named application_dbmail_profile.'.

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that SQL Server Agent jobs running on SQL 1 can notify administrators when a failure occurs.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 52

リモート開発者がDB1とDB2にアクセスできるようにするソリューションを推奨する必要があります。このソリューションは、計画されている変更をサポートし、セキュリティ要件を満たす必要があります。

推薦書には何を含めるべきでしょうか？

- A. データベースレベルのファイアウォールルールを介したパブリックエンドポイント
- B. プライベートエンドポイント
- C. サーバーレベルのファイアウォールルールを介したパブリックエンドポイント
- D. ポイントツーサイト (P2S) VPN

正解: (正解を表示します)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B is Correct: a private endpoint matches the expected DP-300 administration action. A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to recommend a solution that will enable remote developers to access DB1 and DB2.

Why A is Wrong: a public endpoint via a database-level firewall rule is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to recommend a solution that will enable remote developers to access DB1 and DB2.

Why C is Wrong: a public endpoint via a server-level firewall rule is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: a Point-to-Site (P2S) VPN is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to recommend a solution that will enable remote developers to access DB1 and DB2.

質問: 53

あなたは、オンプレミスのアプリケーションをそれぞれAzureに移行するProject1とProject2という2つのプロジェクトに取り組んでいます。Project1は以下の要件を満たす必要があります。

- * アプリをホストするデータベースサーバーのオペレーティングシステムにアクセスできる必要があります。

- * 25TBのストレージ容量が必要です。

プロジェクト2は以下の要件を満たす必要があります。

- * 変更データキャプチャ(CDC)が有効になっていること

- * インスタンスレベルのSQL監査をサポートします。

- * マネージドバックアップをサポートします。

- * 10TBのストレージ容量が必要です。

- * レポートをオフロードします。

プロジェクトに適したAzure SQLソリューションを提案してください。そのソリューションはコストを最小限に抑えるものでなければなりません。

各プロジェクトの推薦書には何を含めるべきでしょうか？回答するには、回答欄で適切な選択肢を選んでください。

注：正解ごとに1ポイントが加算されます。



The screenshot shows the 'Answer Area' for a Microsoft exam. It features a Microsoft logo and two dropdown menus for 'Project1' and 'Project2'. The dropdown menu for 'Project1' is open, showing the following options: Azure Kubernetes Service (AKS), Azure SQL Database, Azure SQL Managed Instance in the Business Critical service tier, Azure SQL Managed Instance in the General Purpose service tier, and SQL Server on Azure Virtual Machines. The dropdown menu for 'Project2' is also open, showing the same options. A large watermark 'jipmiken.com' is visible across the center of the image.

正解:

Answer Area

Project1:

Project2:

Explanation:

Answer Area

Project1:

Project2:

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: Project1 must meet the following requirements: * Have access to the operating system of the database server that hosts the app.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 54

Azure SQL マネージド インスタンス、db1 という名前のデータベース、App1 という名前の Azure Web アプリを含む Azure サブスクリプションがあります。App1 は db1 を使用します。

App1 に対してリソース ガバナーを有効にする必要があります。ソリューションは以下の要件を満たしている必要があります。

App1 は利用可能なすべての CPU リソースを消費できる必要があります。

App1 では、使用可能な CPU リソースの少なくとも半分以上が常に使用可能である必要があります。

どの3つのアクションを順番に実行する必要がありますか？ 回答するには、アクションリストから適切なアクションを解答エリアに移動し、正しい順序に並べます。注：正解の選択肢は複数あります。正解の順序を選択した場合でも、得点は加算されます。

Actions

Create a plan.

Create a classifier function in db1.

Create a workload group.

Create a classifier function in the master database.

Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

正解:

Actions

Create a plan.

Create a classifier function in db1.

Create a workload group.

Create a classifier function in the master database.

Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

Create a workload group.

Create a classifier function in the master database.

Explanation:

Answer Area

Create a resource pool that has the following configurations.

MAX_CPU_PERCENT = 100
MIN_CPU_PERCENT = 50

Create a workload group.

Create a classifier function in the master database.

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to enable Resource Governor for a App1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 55

2つの100GBデータベースをAzureに移動することを計画しています。

ワークロードに基づいてリソース消費を動的にスケーリングする必要があります。

ソリューションは、スケーリング操作中のダウンタイムを最小限に抑える必要があります。

何を使うべきですか？

- A.** AzureSQLデータベースのエラスティックプール
- B.** Azure仮想マシン上のSQL Server
- C.** AzureSQLデータベースのマネージドインスタンス
- D.** AzureSQLデータベース

正解: ([正解を表示します](#))

Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why B is Correct: SQL Server on Azure virtual machines is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. In this scenario, the important constraint is: You need to dynamically scale resources consumption based on workloads. SQL Server on Azure virtual machines satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why C is Wrong: an Azure SQL Database managed instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why D is Wrong: Azure SQL databases is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

質問: 56

R を主要言語としてサポートするが、Scala と SQL もサポートする Azure Databricks で新しいノートブックを作成しています。

言語を切り替えるには、どのスイッチを使用する必要がありますか？

- A.** \[<language>]
- B.** %<language>
- C.** \[<language>]
- D.** @<language>

正解: ([正解を表示します](#))

Explanation/Reference:

You can override the default language by specifying the language magic command % < language > at the beginning of a cell. The supported magic commands are: %python, %r, %scala, and %sql.

Reference:

<https://docs.microsoft.com/en-us/azure/databricks/notebooks/notebooks-use> Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: % < language > is correct because it is the feature whose normal purpose matches the stated requirement. % < language > is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: Which switch should you use to switch between languages?

Why A is Wrong: \[< language >] is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why C is Wrong: \[< language >] is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

Why D is Wrong: @ < language > is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: 57

オンプレミスのMicrosoft SQL Server 2016インスタンスがあり、db1というデータベースをホストしています。Azureサブスクリプションには、MilというAzure SQLマネージドインスタンスが含まれています。


Azure Database Migration Service を使用して、db1 から MM へのオンライン移行を実行する予定です。

移行のためにバックアップを作成する必要があります。ソリューションでは、作成されるバックアップファイルの数を最小限に抑える必要があります。


どのような種類のバックアップを作成し、どのようにバックアップを保存する必要がありますか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area



正解:



Explanation:

Answer Area

Microsoft

Backup type: Full only

To store the backups: Append each backup to a separate file.

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create the backups for the migration.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 58

Server1 という名前の Azure SQL マネージド インスタンスを含む Azure サブスクリプションがあります。

サーバー 1 の監査を有効にする必要があります。ソリューションでは、管理の労力を最小限に抑える必要があります。

どの 4 つのアクションを順番に実行する必要がありますか? 回答するには、アクション リストから適切なアクションを回答領域に移動し、正しい順序に並べます。

Actions

- On Server1, run the create server audit specification T-SQL statement.
- On Server1, run the create server audit T-SQL statement.
- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, enable diagnostics.

Answer Area

Microsoft

正解:

Actions

- On Server1, run the create server audit specification T-SQL statement.
- On Server1, run the create server audit T-SQL statement.
- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, enable diagnostics.

Answer Area

- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, enable diagnostics.

Explanation:

Actions	Answer Area
On Server1, run the create server audit specification T-SQL statement.	1 From the Azure portal, create an Azure Storage account.
On Server1, run the create server audit T-SQL statement.	2 From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
	3 On Server1, create a credential.
	4 On Server1, enable diagnostics.

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to enable auditing for Server 1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 59

新しい Azure SQL マネージド インスタンスができました。

SalesTotal という名前の SQL Server エージェント ジョブを作成する予定です。

SalesTotal が通知を送信できることを確認する必要があります。

コードをどのように完成させるべきですか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area



Microsoft

EXECUTE msdb.dbo.sysmail_add_account_sp

```
@account_name = 'SQL Agent Account',  
@description = 'Mail account',  
@email_address = 'salesdirector@contoso.com',  
@display_name = 'SQL Agent Account',  
@mailserver_name = 'smtp1.contoso.com',
```

EXECUTE msdb.dbo.sysmail_add_profile_sp

```
@profile_name =  
@description = 'E-mail profile'  
'AzureManagedInstance_dbmail_profile',  
'AzureManagedInstanceDefault_profile',  
'AzureManagedInstanceSystemEmail_Profile',
```

EXEC msdb.dbo.sp_add_job @job_name=N'SalesTotals',

@enabled=1,

@notify_level_eventlog=0,

@notify_level_netsend=2,

@notify_level_email=2,

@notify_level_pager=2,

@description=N'Find Sales Totals',

@category_name=N'[Uncategorized (local)]',

@owner_login_name=N'sa'

正解:

Answer Area

```
EXECUTE msdb.dbo.sysmail_add_account_sp
    @account_name = 'SQL Agent Account',
    @description = 'Mail account',
    @email_address = 'salesdirector@contoso.com',
    @display_name = 'SQL Agent Account',
    @mailserver_name = 'smtp1.contoso.com',

EXECUTE msdb.dbo.sysmail_add_profile_sp
    @profile_name = 'AzureManagedInstance_dbmail_profile',
    @description = 'AzureManagedInstanceDefault_profile',
    @notify_level_eventlog=0,
    @notify_level_netsend=2,
    @notify_level_email=2,
    @notify_level_page=2,
    @description=N'Find Sales Totals',
    @category_name=N'[Uncategorized (Local)]',
    @owner_login_name=N'sa'
```



Explanation:



Microsoft

```
EXECUTE msdb.dbo.sysmail_add_account_sp
    @account_name = 'SQL Agent Account',
    @description = 'Mail account',
    @email_address = 'salesdirector@contoso.com',
    @display_name = 'SQL Agent Account',
    @mailserver_name = 'smtp1.contoso.com',

EXECUTE msdb.dbo.sysmail_add_profile_sp
    @profile_name = 'AzureManagedInstanceDefault_profile',
    @description = 'E-mail profile.';

EXEC msdb.dbo.sp_add_job @job_name=N'SalesTotals',
    @enabled=1,
    @notify_level_eventlog=0,
    @notify_level_netsend=2,
    @delete_level=0,
    @description=N'Find Sales Totals',
    @category_name=N'[Uncategorized (local)]',
    @owner_login_name=N'sa'
```

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that SalesTotal can send out notifications.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 60

SQL Server 2019 または Windows Server 2019 を実行する Azure 仮想マシンがあります。仮想マシンには 4 つの vCPU と 28 GB のメモリがあります。

仮想マシンを 8 個の vCPU と 64 GB のメモリにスケールアップします。

サーバーのパフォーマンスに悪影響を与えずに、tempdb の競合を減らす必要があります。

tempdb に構成する必要があるセカンダリ データ ファイルの数はいくつですか？

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

正解: C ([コメントを发表する](#))

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why C is Correct: 8 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to reduce tempdb contention without negatively affecting server performance. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: 2 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: 4 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: 64 is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 61

DB1 という名前の Azure SQL データベースを含む Azure サブスクリプションがあります。

Azure コマンドライン インターフェイス (Azure CLI) を使用して、DB1 を仮想コアベースの購入モデルに移行する必要があります。このソリューションは、以下の要件を満たす必要があります。

* 2 つの高可用性レプリカを含めます。

* ダウンタイムを最小限に抑えます。

Azure CLI コマンドをどのように完了すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

```
az sql db update -g RG1 -s Server1 -n DB1 --edition
```

GeneralPurpose
Hyperscale
Premium
Standard

--service-objective HS_Gen5_2 --read-replicas 2

--manual-cutover
--perform-cutover
--read-scale
--zone-redundant



正解:

Answer Area

```
az sql db update -g RG1 -s Server1 -n DB1 --edition
```

GeneralPurpose
Hyperscale
Premium
Standard

--service-objective HS_Gen5_2 --read-replicas 2

--manual-cutover
--perform-cutover
--read-scale
--zone-redundant



Explanation:

Answer Area

```
az sql db update -g RG1 -s Server1 -n DB1 --edition
```

Hyperscale

--service-objective HS_Gen5_2 --read-replicas 2

--zone-redundant

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to migrate DB1 to a vCore-based purchasing model by using the Azure Command-Line Interface (Azure CLI).

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 62

DB1 という名前のデータベースを含む Azure 仮想マシン上に SQL Server があります。

データベースは CHECKSUM エラーを報告します。

データベースを回復する必要があります。

どのように文を完成させればよいでしょうか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

USE master;

ALTER DATABASE [DB1] SET

▼
OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY

 WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

▼
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST

 WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

▼
MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;

GO

正解:

USE master;

ALTER DATABASE [DB1] SET

▼
OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY

 WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

▼
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST

 WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

▼
MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;

GO



Explanation:

USE master;

ALTER DATABASE [DB1] SET

▼
OFFLINE
ONLINE
SINGLE_USER
TRUSTWORTHY

 WITH ROLLBACK IMMEDIATE;

GO

DBCC CHECKDB ('DB1',

▼
MOINDEX
PHYSICAL_ONLY
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST

 WITH NO_INFOMSGS;

GO

ALTER DATABASE [DB1] SET

▼
MULTI_USER;
ONLINE;
OPEN;
TRUSTWORTHY;

GO



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.

Reference:

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

質問: 63

DB1 という名前のデータベースをホストするオンプレミスの Microsoft SQL Server 2019 インスタンスがあります。

Azure Database Migration Service を使用して、DB1 から Azure SQL マネージド インスタンスへのオンライン移行を実行する予定です。

Azure Database Migration Service からアクセスできる DB1 のバックアップを作成する必要があります。

バックアップでは何を実行し、バックアップはどこに保存すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Run:

A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in:

A Recovery Services vault
An Azure Blob storage account
An SMB file share

正解:

Run:

A full backup and a log backup appended to the same file by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH CHECKSUM option
A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option

Store the backup in:

A Recovery Services vault
An Azure Blob storage account
An SMB file share

Explanation:

The screenshot shows a configuration window for a backup operation. It has two main sections: 'Run:' and 'Store the backup in:'. The 'Run:' section has a dropdown menu with three options: 'A full backup and a log backup appended to the same file by using the WITH CHECKSUM option' (which is highlighted), 'A full backup and a log backup to separate files by using the WITH CHECKSUM option', and 'A full backup and a log backup to separate files by using the WITH FILE_SNAPSHOT option'. The 'Store the backup in:' section has a dropdown menu with three options: 'A Recovery Services vault', 'An Azure Blob storage account' (which is highlighted), and 'An SMB file share'.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-managed-instance-online> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create a backup of DB1 that is accessible to the Azure Database Migration Service.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 64

Microsoft SQL Server インスタンスをホストする Azure 仮想マシンをデプロイするには、Azure Resource Manager (ARM) テンプレートを使用する必要があります。このソリューションでは、SQL Server データベースとログファイルのディスク I/O パフォーマンスを最大化する必要があります。テンプレートはどのように完成させるべきでしょうか？回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

```

"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,
    ...
  }
}

"resources": [
  ...
  "osDisk": {
    ...
    "copy": [
      {
        "name": "dataDisks","count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
        "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",
        "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),parameters('dataDiskCount')),
        variables('dataDisks').caching )]", "diskSizeGB": 1023,

```

Microsoft

None
ReadOnly
ReadWrite

None
ReadOnly
ReadWrite

正解:

```

"variables": {
  "dataDisks": {
    "caching":  "dataDiskCount": 8,"logDisksCount": 1,
    ...
  }
}

"resources": [
  ...
  "osDisk": {
    ...
    "copy": [
      {
        "name": "dataDisks","count": "[add(variables('dataDiskCount'), variables('logDisksCount'))]",
        "input": {"lun": "[copyIndex('dataDisks')]", "createOption": "empty",
        "caching": "[if(greaterOrEquals(copyIndex('dataDisks'),parameters('dataDiskCount')),
        variables('dataDisks').caching )]", "diskSizeGB": 1023,

```

Microsoft

None
ReadOnly
ReadWrite

None
ReadOnly
ReadWrite

Explanation:

Read only
ReadWrite

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to use an Azure Resource Manager ARM) template to deploy an Azure virtual machine that will host a Microsoft SQL Server instance.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 65

論理SQLサーバーを含むAzureサブスクリプションがあります。このサーバーは、db1とdb2という2つのデータベースと、app1というAzure ADサービスプリンシパルをホストしています。app1 が db1 にアクセスできるようにする必要があります。このソリューションでは、最小権限の原則を適用する必要があります。

Transact-SQL ステートメントをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

CREATE

USER

CREDENTIAL

LOGIN

USER

[app1]

FROM EXTERNAL PROVIDER

FOR LOGIN app1

FROM EXTERNAL PROVIDER

FROM LOGIN app1

WITHOUT LOGIN

Answer Area

CREATE

USER

CREDENTIAL

LOGIN

USER

[app1]

FROM EXTERNAL PROVIDER

FOR LOGIN app1

FROM EXTERNAL PROVIDER

FROM LOGIN app1

WITHOUT LOGIN

正解:

Explanation:

Answer Area

CREATE USER [app1] FROM EXTERNAL PROVIDER

<https://learn.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-service-principal-tutorial?view=azuresql>

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that app1 can access db1.

Why the alternate choices are Wrong: The alternate actions fail because Entra authentication must be established at the server/administrator level first, then mapped to contained database users, and finally used by the client connection.

質問: 66

DB1 という Azure SQL データベースがあり、そこには Table1 と Table2 という 2 つのテーブルが含まれています。どちらのテーブルにも Column1 という列が含まれています。Column1 は、App1 というアプリケーションによって結合に使用されます。

保存中、転送中、および使用中の Column1 の内容を保護する必要があります。

列1の内容をどのように保護する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Encryption key:

▼

Column encryption key

Database encryption key

Service master key

Encryption type:

▼

Deterministic

Randomized

Transparent Data Encryption (TDE)

正解:

Encryption key:

▼

Column encryption key

Database encryption key

Service master key

Encryption type:

▼


Deterministic

Randomized

Transparent Data Encryption (TDE)

Explanation:

Encryption key: 

Encryption type: 

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.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine> Basic Concept: This question tests implement a secure environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to protect the contents of Column1 at rest, in transit, and in use.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 67

Microsoft SQL Server 2022 Enterprise がインストールされているオンプレミス サーバーが 5 台あります。各サーバーには、Web アプリのデータを格納し、サードパーティのバックアップ ソリューションを使用してバックアップされる複数のデータベースが含まれています。

データベースを Azure に移行する予定です。

データベースをホストするためのソリューションを推奨する必要があります。ソリューションは以下の要件を満たす必要があります。

* コンピューティング リソースとストレージ リソースは、データベース間で共有する必要があります。

* コストを最小限に抑える必要があります。

推薦書には何を含めるべきでしょうか？

- A. Azure 仮想マシン上の SQL Server
- B. Azure SQL Database エラスティック プール内
- C. Azure SQL データベース
- D. Azure SQL マネージドインスタンス

正解: ([正解を表示します](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. In this scenario, the important constraint is: You have five on-premises servers that have Microsoft SQL Server 2022 Enterprise installed Each server contains multiple database that store data for web apps and are backed up by using a third-party backup solution. in Azure SQL Database elastic pool satisfies that constraint without adding an unrelated service or manual process.

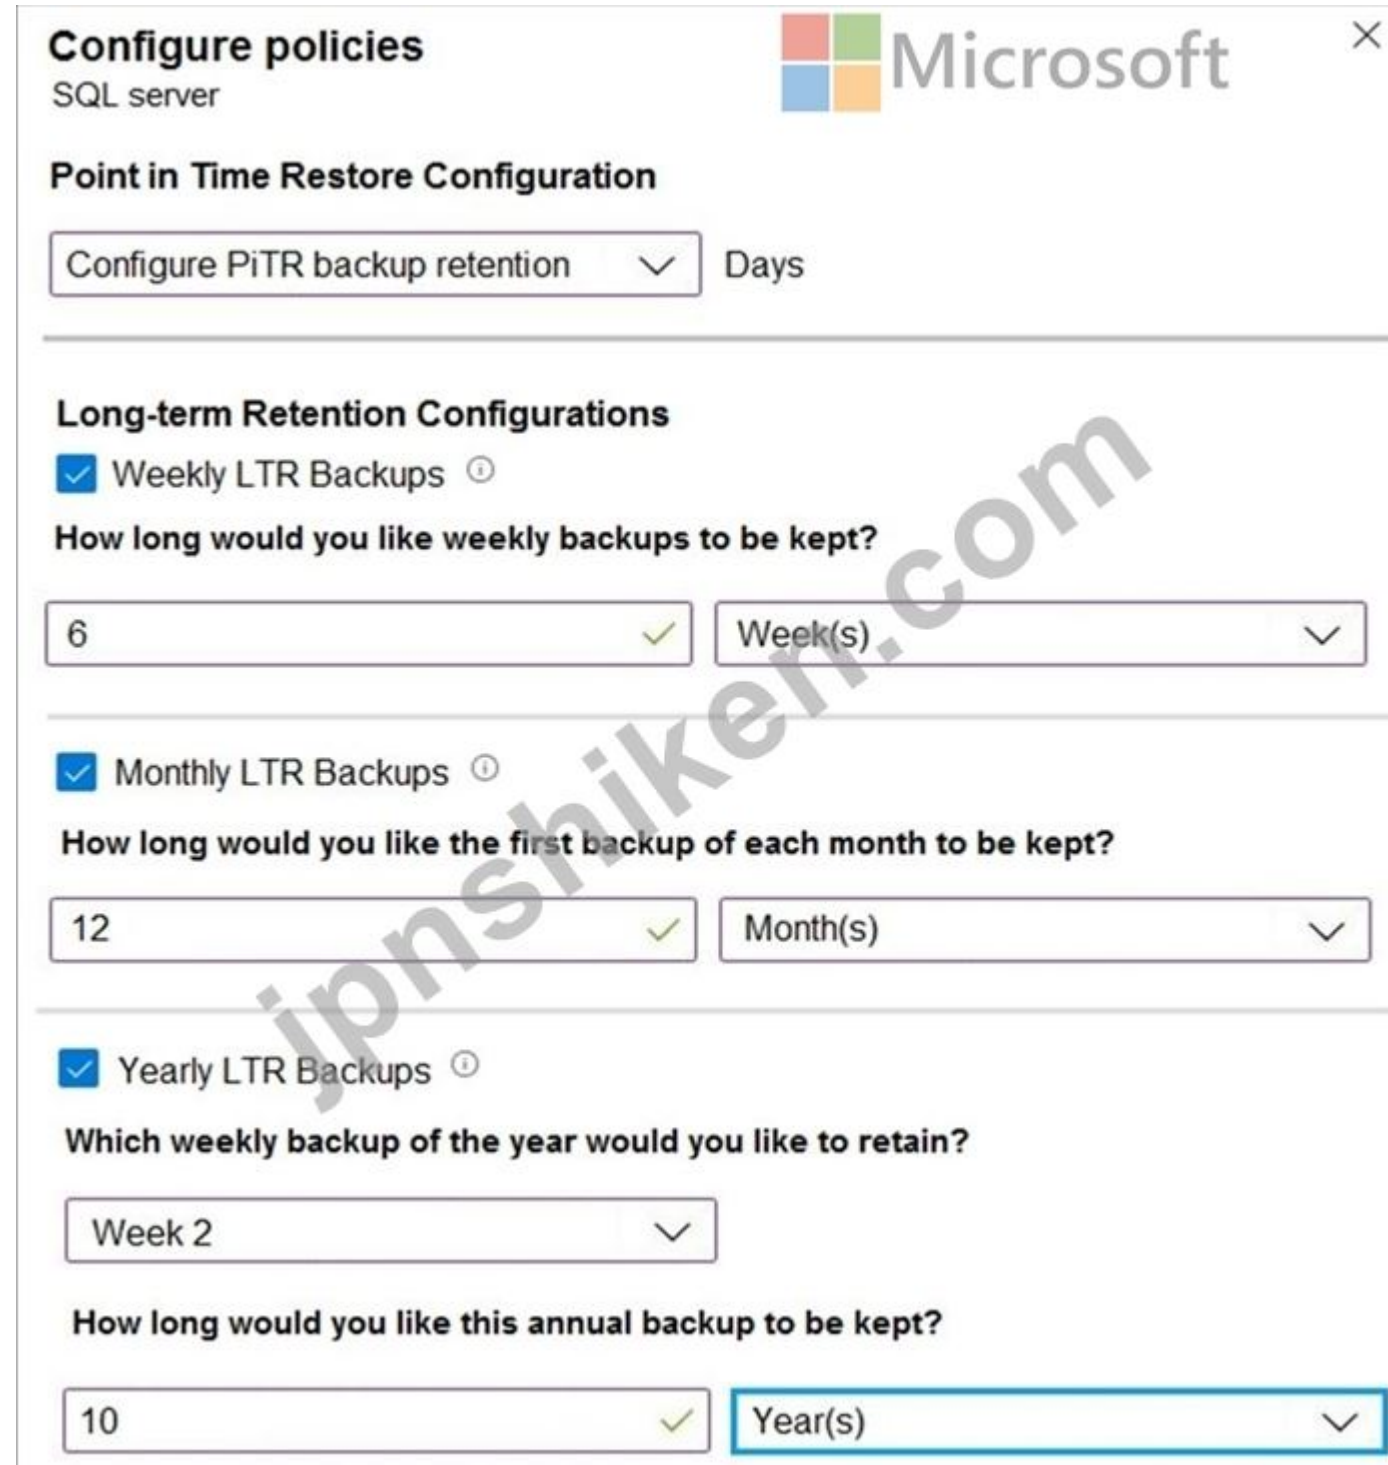
Why A is Wrong: SQL Server on Azure Virtual Machines is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Azure SQL Database is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why D is Wrong: Azure SQL Managed Instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

質問: 68

図に示すように、Azure SQL データベースの長期保有ポリシーを構成します。(図タブをクリックします。)



The screenshot shows the 'Configure policies' dialog for a SQL server. The 'Point in Time Restore Configuration' section has a dropdown set to 'Configure PiTR backup retention' and a unit of 'Days'. The 'Long-term Retention Configurations' section has three checked options: 'Weekly LTR Backups', 'Monthly LTR Backups', and 'Yearly LTR Backups'. For 'Weekly LTR Backups', the retention is set to 6 weeks. For 'Monthly LTR Backups', the retention is set to 12 months. For 'Yearly LTR Backups', the retention is set to 10 years, and the specific backup to retain is 'Week 2'. The 'Year(s)' dropdown is highlighted with a blue border.

Configure policies
SQL server

Point in Time Restore Configuration

Configure PiTR backup retention Days

Long-term Retention Configurations

☒ Weekly LTR Backups

How long would you like weekly backups to be kept?

6 Week(s)

☒ Monthly LTR Backups

How long would you like the first backup of each month to be kept?

12 Month(s)

☒ Yearly LTR Backups

Which weekly backup of the year would you like to retain?

Week 2

How long would you like this annual backup to be kept?

10 Year(s)

最初の週次バックアップは 2020 年 1 月 4 日に実行されました。最初の 10 週次バックアップの日付は次のとおりです。

- * 2020年1月4日
- * 2020年1月11日
- * 2020年1月18日
- * 2020年1月25日
- * 2020年2月1日
- * 2020年2月8日
- * 2020年2月15日
- * 2020年2月22日
- * 2020年2月29日
- * 2020年3月7日

ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。

注意: 正しい選択ごとに 1 ポイントが加算されます。




The backup saved to long-term retention on January 4, 2020, will be retained for

The backup saved to long-term retention on January 11, 2020 will be retained for

Options for both questions: 6 weeks, 12 months, 10 years

正解:



The backup saved to long-term retention on January 4, 2020, will be retained for

The backup saved to long-term retention on January 11, 2020 will be retained for

Options for both questions: 6 weeks, 12 months, 10 years

Explanation:

The backup saved to long-term retention on January 4, 2020, will be retained for

	▼
6 weeks	
12 months	
10 years	

The backup saved to long-term retention on January 11, 2020 will be retained for

	▼
6 weeks	
12 months	
10 years	



Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You configure a long-term retention policy for an Azure SQL database as shown in the exhibit.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 69

databasebackups という名前のストレージ アカウントを含む Azure サブスクリプションがあります。

DB1 という名前の Azure SQL マネージド インスタンスがあります。

DB1 をデータベースバックアップにバックアップする必要があります。

コマンドをどのように完了すればよいですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

CREATE CREDENTIAL

[https://databasebackups.blob.core.windows.net/Backups]

WITH IDENTITY =

'SHARED ACCESS SIGNATURE'	▼
'DatabaseBackups'	
'KeyVault1'	
'SHARED ACCESS SIGNATURE'	

SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-

02T19:30:08Z&spr=https&sv=2021-06-

08&sr=b&sig=B%2FxEYQiOC%4BqyYCeqlWShz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'

BACKUP DATABASE DB1

TO URL =

'https://databasebackups.blob.core.windows.net/Backups/db1.bak'

WITH

COPY_ONLY	▼
CHECKSUM	
COMPRESSION	
COPY_ONLY	
DIFFERENTIAL	

正解:

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE'
SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-
02T19:30:08Z&spr=https&sv=2021-06-
08&sr=b&sig=B%2FxEYQ10C%4BqyYCeQWHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL =
'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY_ONLY
```

Explanation:

Answer Area

```
CREATE CREDENTIAL
[https://databasebackups.blob.core.windows.net/Backups]
WITH IDENTITY = 'SHARED ACCESS SIGNATURE' ,
SECRET = 'sp=r&st=2023-02-02T19:23:08Z&se=2033-02-
02T19:30:08Z&spr=https&sv=2021-06-
08&sr=b&sig=B%2FxEYQ10C%4BqyYCeQWHSz2QpRI%2FKcg3ZABz78J2kix3JZjk%3D'
BACKUP DATABASE DB1
TO URL =
'https://databasebackups.blob.core.windows.net/Backups/db1.bak'
WITH COPY_ONLY
```

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to back up DB1 to databasebackups.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 70

Azure SQL マネージド インスタンスと、バックアップを実行するジョブを作成します。

ジョブが失敗した場合に、配布グループに電子メールで通知するようにジョブを設定する必要があります。このソリューションは、管理作業を最小限に抑えるものでなければなりません。

どの3つの行動を順番に実行すべきでしょうか？回答するには、行動リストから適切な行動を回答欄に移動させ、正しい順序に並べ替えてください。

注：正解の選択肢は複数あります。どの正解の選択肢を選んでも点数が加算されます。

Actions Commands Cmdlets Statements

- Configure SendGrid.
- Configure an alert.
- Configure Database Mail.
- Create an operator.
- Configure a job notification.

Answer Area

正解:

Actions Commands Cmdlets Statements

- Configure SendGrid.
- Configure an alert.
- Configure Database Mail.
- Create an operator.
- Configure a job notification.

Answer Area

- Configure Database Mail.
- Create an operator.
- Configure a job notification.

Explanation:

Actions Commands Cmdlets Statements

⋮ Configure SendGrid.
⋮ Configure an alert.

Answer Area

1	⋮ Configure Database Mail.
2	⋮ Create an operator.
3	⋮ Configure a job notification.



Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to configure the job to notify a distribution group by email when the job fails.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 71

20 個の Azure SQL データベースを含む Azure サブスクリプションがあります。

データベースのインデックス メンテナンスを実行するには、Transact-SQL ステートメントを作成します。

Transact-SQL コマンドを使用して、各データベースに対して 1 日 1 回ステートメントを実行するようにスケジュールする必要があります。

ステートメントをスケジュールするには何を使用すればよいですか？

- A. Azure オートメーション
- B. 弾力性のある仕事
- C. SQL Server エージェント ジョブ
- D. Azure関数

正解: ([正解を表示します](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: Azure Automation runbooks execute PowerShell or Python tasks on a schedule, making them useful for Azure SQL maintenance workflows. This fits the case because the requested outcome is: You need to schedule the statement to run once daily against each database by using Transact-SQL commands.

The selected feature addresses that outcome directly rather than relying on a workaround.

Why B is Wrong: Elastic Jobs automate T-SQL execution across many Azure SQL databases, especially for multi-tenant or shard-style deployments. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

Why C is Wrong: SQL Server Agent schedules and runs SQL Server jobs; in Azure SQL Database, Elastic Jobs or Azure Automation are often used instead. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why D is Wrong: an Azure function can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

質問: 72

DB3の管理を自動化するプロセスを推奨する必要があります。ソリューションは管理要件を満たす必要があります。プロセスの最初のステップは何でしょうか？

- A. DB3 をホストする論理サーバーの Microsoft Entra 認証を構成します。

- B. データベース スコープの資格情報を持つデータベースを作成します。
- C. DB3 への接続用のプライベート エンドポイントを構成します。
- D. DB3 にデータベース スコープの資格情報を作成します。

正解: ([正解を表示します](#))

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why C is Correct: A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. The scenario asks for: You need to recommend a process to automate the management of DB3. That makes Configure a private endpoint for connectivity to DB3. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Configure Microsoft Entra authentication for the logical server that hosts DB3. is a security- related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why B is Wrong: Create a database that has database-scoped credentials. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to recommend a process to automate the management of DB3.

Why D is Wrong: Create data base-scoped credentials in DB3. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

質問: 73

VM1という名前のカスタムイメージに基づくAzure仮想マシンがあります。

VM1は、Microsoft SQL Server 2019Standardのインスタンスをホストします。

次の要件を満たすには、VM1のメンテナンスを自動化する必要があります。

SQLServerとWindowsServerのパッチを自動化します。

VM1上のデータベースの完全なデータベースバックアップとトランザクションログバックアップを自動化します。

管理作業を最小限に抑えます。

あなたは最初に何をすべきですか？

- A. VM1に対してシステムによって割り当てられたマネージIDを有効にします
- B. VM1をMicrosoft.Sqlリソースプロバイダーに登録します
- C. Azure仮想マシンのDesired State Configuration (DSC) 拡張機能をVM1にインストールします
- D. VM1をMicrosoft.SqlVirtualMachineリソースプロバイダーに登録します

正解: ([正解を表示します](#))

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 Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: Register VM1 to the Microsoft.SqlVirtualMachine resource provider is correct because it is the feature whose normal purpose matches the stated requirement. Register VM1 to the Microsoft.

SqlVirtualMachine resource provider is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario wording points to that specific behavior: You need to automate the maintenance of VM1 to meet the following requirements: Automate the patching of SQL Server and Windows Server.

Why A is Wrong: A managed identity gives an Azure resource an identity in Microsoft Entra ID so it can access services such as Key Vault without embedding credentials in code. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why B is Wrong: Register VM1 to the Microsoft.Sql resource provider is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why C is Wrong: Install an Azure virtual machine Desired State Configuration (DSC) extension on VM1 is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

質問: 74

同じ論理サーバー上に、Azure SQL Databaseのインスタンスを5つ作成します。

各データベースで、User1 という名前の Azure Active Directory (Azure AD) ユーザーのユーザーを作成します。

ユーザー1はAzure Data Studioを使用して論理サーバーへの接続を試みたが、ログインエラーが発生した。

User1がAzure Data Studioを使用して論理サーバーに接続した際に、User1がすべてのデータベースを表示できることを確認する必要があります。

あなたはどうすべきでしょうか？

A. マスターデータベースにUser1を作成します。

B. master データベースの sys.databases に対する select 権限を public に付与します。

C. ユーザー1にマスターデータベースのdb_datareaderロールを割り当てます。

D. User1 が作成するデータベースに対して、User1 に db_datareader ロールを割り当てます。

正解: A ([コメントを發表する](#))

質問: 75

PaaS ソリューションの災害復旧要件を満たすには何を実装する必要がありますか？

A. 可用性ゾーン

B. フェイルオーバーグループ

C. Always On 可用性グループ

D. 地理的レプリケーション

正解: ([正解を表示します](#))

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 geo-replication 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/automatic-failover-group-overview>

質問: 76

SQLMI1という名前のAzureSQLDatabaseマネージドインスタンスがあります。Microsoft SQLServerエージェントジョブはSQLMI1で実行されます。

ジョブが完了したら、自動電子メール通知が送信されるようにする必要があります。

ソリューションに何を含める必要がありますか？

A. SQL Server構成マネージャー (SSMS) から、SQLServerエージェントを有効にします

B. SQL Server Management Studio (SSMS) から、sp_set_sqlagent_propertiesを実行します

C. SQL Server Management Studio (SSMS) から、データベースメールプロファイルを作成します

D. Azureポータルから、Eメール/ SMS /プッシュ/音声アクションを持つAzureMonitorアクショングループを作成します

正解: **C** ([コメントを发表する](#))

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:

In Object Explorer, expand a SQL Server instance.

Right-click SQL Server Agent, and then click Properties.

Click Alert System.

Select Enable Mail Profile.

In the Mail system list, select Database Mail.

In the Mail profile list, select a mail profile for Database Mail.

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>

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: **77**

タスク3

dbl に対して実行されたすべてのクエリがクエリ ストアにキャプチャされるようにする必要があります。

正解:

See the explanation part for the complete Solution.

Explanation:

To ensure that all queries executed against dbl are captured in the Query Store, you need to enable the Query Store feature for the database and set the query capture mode to ALL. The Query Store feature provides you with insight on query plan choice and performance for Azure SQL Database1. The query capture mode controls whether all queries or only a subset of queries are tracked2.

Here are the steps to enable the Query Store and set the query capture mode to ALL for the database dbl:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server.

Select the database dbl and click on Query Performance Insight in the left menu.

Click on Configure Query Store and turn on the Query Store switch.

In the Query Capture Mode dropdown, select All and click on Save.

Using Transact-SQL statements:

Connect to the Azure SQL Database server and the database dbl using SQL Server Management Studio or Azure Data Studio.

Run the following command to enable the Query Store for the database: ALTER DATABASE db1 SET QUERY_STORE = ON; Run the following command to set the query capture mode to ALL for the database: ALTER DATABASE db1 SET QUERY_STORE (QUERY_CAPTURE_MODE = ALL); These are the steps to ensure that all queries executed against db1 are captured in the Query Store.

質問: 78

Server1 というサーバーに DatabaseA という Azure SQL Database インスタンスがあります。
App1 という新しいユーザーを DatabaseA に追加し、App1 に db_datacenter 権限を付与する予定です。App1 は SQL Server 認証を使用します。
App1 を作成する必要があります。ソリューションでは、App1 が同じ資格情報を使用して他のデータベースにアクセスできるようにする必要があります。
順番に実行する必要がある 3 つのアクションはどれですか。回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions

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];

Answer Area

Microsoft

正解:

Actions

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];

Answer Area

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];

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/>

質問: 79

DB1 という名前の Azure SQL データベースを含む Azure サブスクリプションがあります。

DB1のクエリストアキャプチャモードを構成する必要があります。ソリューションは以下の要件を満たす必要があります。

* リソース消費量が大きいクエリのみが追跡されるようにします。

* 管理作業を最小限に抑える

クエリ ストア キャプチャ モードを何に設定すればよいですか?

- A. なし
- B. カスタム
- C. 自動
- D. すべて

正解: (正解を表示します)

Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: Auto matches the expected DP-300 administration action. Auto is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to configure Query Store Capture Mode for DB1.

Why A is Wrong: None is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why B is Wrong: Custom is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: All is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 80

D61 という名前の Azure SQL データベースがあります。

DB1 に割り当てられた未使用スペースの量(メガバイト単位)を特定する必要があります。

Transact-SQL クエリをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

```
SELECT SUM(size/ 128.0) - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0 AS DatabaseDataSpaceAllocatedUnusedInMB
FROM sys.databases
GROUP BY sys.database_files
HAVING t.sys.dm_db_resource_stats
```

正解:

Answer Area

```
SELECT SUM(size/ 128.0) - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0 AS DatabaseDataSpaceAllocatedUnusedInMB
FROM sys.databases
GROUP BY sys.database_files
HAVING t.sys.dm_db_resource_stats
```

Explanation:

Answer Area

```
SELECT SUM(size/ 128.0) - CAST(FILEPROPERTY(name, 'SpaceUsed') AS int)/128.0 AS DatabaseDataSpaceAllocatedUnusedInMB
FROM sys.database_files
GROUP BY type_desc
HAVING type_desc = 'ROWS';
```


Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to identify how much unused space in megabytes was allocated to DB1.


Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 81

PipelineA と PipelineB という名前の 2 つのパイプラインを持つ Azure データ ファクトリがあります。
PipelineA には、次の図に示すように 4 つのアクティビティがあります。



PipelineB には、次の図に示すように 2 つのアクティビティがあります。



データファクトリのアラートを作成し、パイプラインとすべての失敗の種類の両方について 失敗したパイプライン実行」メトリックを使用します。メトリックの設定は次のとおりです。

- * 演算子: より大きい
- * 集計タイプ: 合計
- * 閾値: 2
- * 集計粒度(期間): 5分
- * 評価頻度: 5分ごと

Data Factory の監視では、次の表に示す障害が記録されます。

Pipeline	Activity	Time
PipelineA	Activity1	31-Jan-2020 10:44:00
PipelineA	Activity3	31-Jan-2020 10:47:00
PipelineB	Activity1	31-Jan-2020 10:50:00

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	☐	☐
An alert notification was sent after the failure of Activity3 in PipelineA.	☐	☐
An alert notification was sent after the failure of Activity1 in PipelineB.	☐	☐

正解:

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	☐	☒
An alert notification was sent after the failure of Activity3 in PipelineA.	☐	☒
An alert notification was sent after the failure of Activity1 in PipelineB.	☐	☒

Explanation:

Statements	Yes	No
An alert notification was sent after the failure of Activity1 in PipelineA.	☐	☒
An alert notification was sent after the failure of Activity3 in PipelineA.	☐	☒
An alert notification was sent after the failure of Activity1 in PipelineB.	☐	☒

Box 1: No

Just one failure within the 5-minute interval.

Box 2: No

Just two failures within the 5-minute interval.

Box 3: No

Just two failures within the 5-minute interval.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-overview> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure data factory that has two pipelines named PipelineA and PipelineB.

Why the alternate choices are Wrong: The alternate destinations or collection scopes miss part of the telemetry requirement: they either omit the server/pool/database layer or send data somewhere that does not support the requested analysis path.

質問: 82

Azure Synapse Analyticsで、Customersという名前のテーブルを含むエンタープライズデータウェアハウスを設計しています。顧客にはクレジットカード情報が含まれます。

営業担当者に顧客のすべてのエントリを表示する機能を提供するソリューションを推奨する必要があります。

このソリューションでは、すべての営業担当者がクレジットカード情報を表示または推測できないようにする必要があります。

推奨事項には何を含める必要がありますか？

- A. 行レベルのセキュリティ
- B. データマスキング
- C. 常に暗号化
- D. 列レベルのセキュリティ

正解: ([正解を表示します](#))

Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics support dynamic data masking. Dynamic data masking limits sensitive data exposure by masking it to non-privileged users.

The Credit card masking method exposes the last four digits of the designated fields and adds a constant string as a prefix in the form of a credit card.

Example:

XXXX-XXXX-XXXX-1234

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B is Correct: data masking matches the expected DP-300 administration action. data masking is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to recommend a solution to provide salespeople with the ability to view all the entries in Customers.

Why A is Wrong: Row-Level Security filters rows returned by queries based on a predicate function, allowing different users or tenants to see different subsets of the same tables. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why C is Wrong: Always Encrypted protects sensitive columns from exposure to the database engine by encrypting and decrypting data at the client side with keys protected outside the database. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why D is Wrong: column-level security is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

質問: **83**

DB1 という名前の Azure SQL データベースがあります。

データセンターに障害が発生した場合でも、DB1がデータ損失のない自動フェイルオーバーをサポートできるようにする必要があります。ソリューションはコストを最小限に抑える必要があります。どの展開オプションと価格帯を構成する必要がありますか？

- A. Azure SQL データベース ハイパースケール
- B. Azure SQL Database マネージドインスタンス 汎用
- C. Azure SQL データベース プレミアム
- D. Azure SQL データベース ベーシック

正解: **C** ([コメントを发表する](#))

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.

Reference:

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

質問: **84**

次のリソースを含む Azure サブスクリプションがあります。

- * 10 個の Azure SQL データベース
- * 5 つの Azure SQL マネージド インスタンス
- * Azure Virtual Machines 上の SQL Server の 5 つのインスタンス

すべての Azure SQL リソースに対して集中監視ソリューションを実装する必要があります。ソリューションでは、管理労力を最小限に抑える必要があります。ソリューションには何を含めるべきでしょうか？

- A. ログ分析
- B. Azure SQL アナリティクス

C. クエリ パフォーマンス インサイト

D. SQL インサイト

正解: ([正解を表示します](#))

Basic Concept: This question tests identity and authentication for Azure SQL and SQL Server workloads, including when to use contained users, directory identities, certificate validation, or authentication profiles.

Why B is Correct: Azure SQL Analytics is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You have an Azure subscription that contains the following resources: * 10 Azure SQL databases * Five Azure SQL managed instances * Five instances of SQL Server on Azure Virtual Machines You need to implement a centralized monitoring solution for all the Azure The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Log Analytics stores logs and metrics in a queryable workspace and supports workbooks, dashboards, alerting, and cross-resource analysis. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Query Performance Insight surfaces resource-consuming queries and database workload patterns in Azure SQL for performance investigation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: SQL Insights is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: 85

SVR1からInstance1にデータベースを移行する際に使用するサービスとターゲットエンドポイントを推奨する必要があります。ソリューションは可用性要件を満たす必要があります。

何を推奨しますか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:



Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to recommend which service and target endpoint to use when migrating the databases from SVR1 to Instance1.

Why the alternate choices are Wrong: The alternate choices do not satisfy the stated availability design because quorum and listener resources must match Azure failover-cluster behavior and the required SLA.

質問: 86


Azure Synapse Analytics で、Web サイトのトラフィック分析をスター スキーマに保存するエンタープライズ データ ウェアハウスを設計しています。

ウェブサイト訪問に関するファクトテーブルを作成する予定です。テーブルのサイズは約5GBになります。


テーブルに使用する分散タイプとインデックスタイプを推奨する必要があります。ソリューションは、最速のクエリパフォーマンスを提供する必要があります。

何を推奨しますか? 回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:

Distribution:  ▼

- Hash
- Round robin
- Replicated

Index: ▼

- Clustered columnstore
- Clustered
- Nonclustered

Box 1: Hash

Consider using a hash-distributed table when:

The table size on disk is more than 2 GB.

The table has frequent insert, update, and delete operations.

Box 2: Clustered columnstore

Clustered columnstore tables offer both the highest level of data compression and the best overall query performance.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-distribute>

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-index> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to recommend which distribution type and index type to use for the table.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 87

DB1という名前のAzureSQLデータベースがあります。

タイプCHAR (400)の列が20個あるテーブル名Table1があります。Table1の行圧縮が有効になっています。

データベース監査中に、150文字を超えるフィールドが含まれていないことがわかりました。

Table1にページ圧縮を適用できることを確認する必要があります。

あなたは何をすべきか？

- A. 列をスパースとして構成します。
- B. 列タイプをnvarchar (MAX)に変更します。
- C. 列タイプをvarchar (MAX)に変更します。
- D. 列タイプをvarchar (200)に変更します。

正解: ([正解を表示します](#))

Reference:

<https://www.sqlshack.com/sql-varchar-data-type-deep-dive/>

<https://36chambers.wordpress.com/2020/06/18/nvarchar-everywhere-a-thought-experiment/> Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why D is Correct: Change the column type to varchar (200). is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario.

The scenario asks for: You need to ensure that you can apply page compression to Table1. That makes Change the column type to varchar (200). the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Configure the columns as sparse. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why B is Wrong: Change the column type to nvarchar (MAX). is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Change the column type to varchar (MAX). is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

質問: 88

Azure SQL データベースを含む Azure サブスクリプションがあります。

データベースがクエリにタイムリーに応答できません。

問題が resource_semaphore の待機に関連しているかどうかを特定する必要があります。

Transact-SQL クエリをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

```
SELECT
    is_user_process
    wait_time
    wait_type
SUM(wait_time) AS total_wait_time_ms
FROM sys.
JOIN sys.dm_exec_sessions AS dmvs2
    ON dmvs1.session_id = dmvs2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

正解:

```
SELECT
    is_user_process
    wait_time
    wait_type
    SUM(wait_time) AS total_wait_time_ms
FROM sys.
    dm_exec_query_stats
    dm_exec_requests
    query_store_query
JOIN sys.dm_exec_sessions AS dmv2
    ON dmv1.session_id = dmv2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

Explanation:

```
SELECT
    is_user_process
    wait_time
    wait_type
    SUM(wait_time) AS total_wait_time_ms
FROM sys.
    dm_exec_query_stats
    dm_exec_requests
    query_store_query
JOIN sys.dm_exec_sessions AS dmv2
    ON dmv1.session_id = dmv2.session_id
WHERE is_user_process = 1
GROUP BY wait_type
ORDER BY SUM(wait_time) DESC;
```

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/monitoring-with-dmvs> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to identify whether the issue relates to resource_semaphore waits.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: **89**

タスク2

db1 Azure SQLデータベースの差分バックアップを、1日2回ではなく1日1回に設定する必要があります。SQL Server Management StudioとAzureポータルを使用する必要があるかもしれません。

正解:

See the explanation part for the complete Solution.

Explanation:

Requirement: Configure differential backups for Azure SQL Database db1 to run once a day instead of twice a day.

Correct setting: Change Differential backup frequency from 12 Hours to 24 Hours.

Azure SQL Database supports differential backup frequency of either 12 hours or 24 hours. A 12-hour frequency means twice per day; a 24-hour frequency means once per day. Microsoft also notes that 24-hour differential backup frequency can increase restore time compared with 12-hour frequency.

Method 1 - Azure Portal Method

This is the best method for the simulation.

Step 1: Open the Azure SQL logical server

Sign in to the Azure portal.

Search for SQL servers.

Open the logical SQL server that hosts database db1.

Do not start from SQL Server Management Studio for this task. The differential backup frequency is an Azure SQL backup policy setting, not a normal T-SQL database setting.

Step 2: Open the Backups page

In the SQL server left menu, select Backups.

Select the Retention policies tab.

Microsoft's documented portal path is to go to the logical SQL server, select Backups, then select the Retention policies tab.

Step 3: Select database db1

In the list of databases, locate db1.

Select the checkbox next to db1.

Select Configure policies from the action bar.

Step 4: Change the differential backup frequency

In the policy configuration pane:

Find Differential backup frequency.

Change it from:

12 Hours

to:

24 Hours

Leave the PITR retention period unchanged unless the task specifically tells you to change it.

Select Apply or Save.

Microsoft's documented option is exactly 12 Hours or 24 hours under Differential backup frequency.

Method 2 - Azure CLI Method

Use this if the simulation provides Cloud Shell.

az sql db str-policy set \


```
--resource-group < resource-group-name > \  
--server < server-name > \  
--name db1 \  
--retention-days < current-retention-days > \  
--diffbackup-hours 24
```

Example:

```
az sql db str-policy set \  
--resource-group RG1 \  
--server sql60152867 \  
--name db1 \  
--retention-days 7 \  
--diffbackup-hours 24
```

Microsoft documents `az sql db str-policy set` with `--diffbackup-hours 24` for changing active database differential backup frequency. Valid values are 12 or 24 hours.

Be careful: do not guess the retention days blindly in a real environment. In the exam lab, use the existing retention value shown in the portal unless the task also asks you to change retention.

Method 3 - PowerShell Method

Use this if Azure PowerShell is available.

```
Set-AzSqlDatabaseBackupShortTermRetentionPolicy `  
-ResourceGroupName " < resource-group-name > " `  
-ServerName " < server-name > " `  
-DatabaseName " db1 " `  
-RetentionDays < current-retention-days > `  
-DiffBackupIntervallnHours 24
```

Microsoft documents `Set-AzSqlDatabaseBackupShortTermRetentionPolicy` with `-DiffBackupIntervallnHours 24` for setting Azure SQL Database differential backup frequency.

SSMS / T-SQL Clarification

For this task, SSMS is not the right tool to change the setting.

There is no normal `ALTER DATABASE T-SQL` command in Azure SQL Database to change automated differential backup frequency. Microsoft documents this change through:

Azure portal

Azure CLI

PowerShell

REST API

not SSMS/T-SQL.

You may use SSMS only to confirm the database exists and is accessible, but the backup frequency setting must be changed from Azure management tools.

質問: 90

データベースをAzureSQLDatabaseマネージドインスタンスにデプロイします。

読み取りクエリがデータベースに書き込もうとしているクエリをブロックしないようにする必要があります。

どのデータベースオプションを設定する必要がありますか？

A. 強制へのパラメータ化

B. 単純化へのパラメータ化

C. 強制に対する耐久性の遅延

D. READ_COMMITTED_SNAPSHOTからON

正解: ([正解を表示します](#))

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.

Reference:

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

質問: 91

オンプレミス ネットワークには、DB 1 という名前の 60 TB のデータベースをホストするサーバーが含まれています。ネットワークには 10 Mbps のインターネット接続があります。

DB 1 を Azure に移行する必要があります。ソリューションでは、データベースの移行にかかる時間を最小限に抑える必要があります。

何を使うべきでしょうか？

A. Azure Migrate

B. データ移行アシスタント (DMA)

C. Azure データ ボックス

D. Azure データ ベース移行サービス

正解: C ([コメントを發表する](#))

<https://www.techtarget.com/searchitoperations/tip/Easily-transfer-VMs-to-the-cloud-with-Microsoft-Azure-Migrate-Basic-Concept>: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why C is Correct: Azure Data BOX matches the expected DP-300 administration action. Azure Data BOX is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to migrate DB 1 to Azure.

Why A is Wrong: Azure Migrate is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why B is Wrong: Data Migration Assistant (DMA) is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why D is Wrong: Azure Database Migration Service is the managed migration service used to move database engines to Azure targets with online or offline migration modes depending on source and target support. It applies to a different Azure data-platform design than the one the question is testing.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 92

DB1 という名前の Microsoft SQL Server 2019 データベースと、SQLMI1 という名前の Azure SQL マネージド インスタンスがあります。SQL Server エージェント ジョブを DB1 から SQLMI1 に移動する必要があります。SQLMI1 でサポートされていないジョブ属性はどれですか？

- A. テーブルにログ
- B. メール通知
- C. スケジュール
- D. 出力ファイル

正解: (正解を表示します)

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why D is Correct: output files matches the expected DP-300 administration action. output files can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to move a SQL Server Agent job from DB1 to SQLMI1.

Why A is Wrong: log to table can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why B is Wrong: email notifications can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

Why C is Wrong: schedules can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

質問: 93

AzureSQLデータベースがあります。

プランキャッシュが、一度だけ使用されたコンパイル済みプランでいっぱいであることがわかります。

select * from sys.database_scoped_configurations Transact-SQLコマンドを実行すると、次の表に示す結果が得られます。

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

あなたは記憶の圧力を和らげる必要があります。

何を構成する必要がありますか？

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

正解: C (コメントを发表する)

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.

Reference:

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

質問: 94

RG1とRG2という2つのリソースグループがあります。RG1には、DB1という名前のデータベースをホストするMI1という名前のAzure SQLマネージドインスタンスが含まれています。RG2には、MI2という名前のSQLマネージドインスタンスが含まれています。
以下のPowerShellスクリプトを実行します。

```
$database = 'DB1'
$sourceMI = 'MI1'
$sourceRG = 'RG1'
$targetMI = 'MI2'
$targetRG = 'RG2'

## Start database operation.
$databaseCopyParams = @{
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
}
Copy-AzSqlInstanceDatabase @databaseCopyParams

## Verify the operation status is succeeded.
}
Get-AzSqlInstanceDatabaseCopyOperation @databaseCopyOperationsParams

## Complete database operation.
$completeDatabaseCopyParams = @{
    DatabaseName = $database
    InstanceName = $sourceMI
    ResourceGroupName = $sourceRG
    TargetInstanceName = $targetMI
    TargetResourceGroupName = $targetRG
}
Complete-AzSqlInstanceDatabaseCopy @completeDatabaseCopyParams
```

Answer Area



The copy of DB1 will be dropped from MI1.

A copy of DB1 on MI1 will be available during the operation.

A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.

Yes

No

☐☐☐☐☐☐

正解:

Answer Area

Microsoft

Statements

The copy of DB1 will be dropped from MI1.	Yes ☐	No ☒
A copy of DB1 on MI1 will be available during the operation.	☒	☐
A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.	☐	☒

Explanation:

Answer Area

Microsoft

Statements

The copy of DB1 will be dropped from MI1.	Yes ☐	No ☒
A copy of DB1 on MI1 will be available during the operation.	☒	☐
A copy of DB1 will be restored to MI2 by using a point-in-time restore (PITR) from a backup taken at the time the operation started.	☐	☒

Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have two resource groups named RG1 and RG2.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 95

Azure SynapseAnalytics専用のSQLプールでディメンションテーブルを設計しています。

テーブルの代理キーを作成する必要があります。このソリューションは、最速のクエリパフォーマンスを提供する必要があります。

代理キーには何を使用する必要がありますか？

A. an IDENTITY column

B. a GUID column

C. a sequence object

正解: (正解を表示します)

Dedicated SQL pool supports many, but not all, of the table features offered by other databases.

Surrogate keys are not supported. Implement it with an Identity column.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tablesoverview> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: an IDENTITY column is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to create a surrogate key for the table. The selected feature addresses that outcome directly rather than relying on a workaround.

Why B is Wrong: a GUID column is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: a sequence object is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: **96**

タスク5

db1 にディザスタリカバリソリューションを構成する必要があります。フェイルオーバーが発生した場合、データベースへの接続文字列は変更されません。セカンダリサーバーは、Azure リージョン 米国西部 3」に配置する必要があります。

正解:

See the explanation part for the complete Solution.

Explanation:

To configure a disaster recovery solution for db1, you can use the failover groups feature of Azure SQL Database. Failover groups allow you to manage the replication and failover of a group of databases across different regions with the same connection strings¹. You can also use active geo-replication as an alternative, but you will need to update the connection strings manually after a failover².

Here are the steps to create a failover group for db1 with the secondary server in the West US 3 region:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

Select Failover groups in the left menu and click on Add group.

Enter a name for the failover group and select West US 3 as the secondary region.

Click on Create a new server and enter the details for the secondary server, such as server name, admin login, password, and subscription.

Click on Select existing database(s) and choose db1 from the list of databases on the primary server.

Click on Configure failover policy and select the failover mode, grace period, and read-write failover endpoint mode according to your preferences.

Click on Create to create the failover group and start the replication of db1 to the secondary server.

Using PowerShell commands:

Install the Azure PowerShell module and log in with your Azure account.

Run the following command to create a new server in the West US 3 region: `New-AzSqlServer - ResourceGroupName < your-resource-group-name > -ServerName < your-secondary-server-name > - Location " West US 3 " -SqlAdministratorCredentials $(New-Object -TypeName System.Management.`

`Automation.PSCredential -ArgumentList " < your-admin-login > " , $(ConvertTo-SecureString -String " < your-password > " -AsPlainText -Force))` Run the following command to create a new failover group with db1: `New-AzSqlDatabaseFailoverGroup - ResourceGroupName < your-resource-group-name > -ServerName < your-primary-server-name > - PartnerResourceGroupName < your-resource-group-name > -PartnerServerName < your-secondary-server- name > -FailoverGroupName < your-failover-group-name > -Database db1 -FailoverPolicy Manual -`


`GracePeriodWithDataLossHours 1 -ReadWriteFailoverEndpoint " Enabled " You can modify the parameters of the command according to your preferences, such as the failover policy, grace period, and read-write failover endpoint mode.`

These are the steps to create a failover group for db1 with the secondary server in the West US 3 region.

質問: **97**

Azure SQL データベースがあります。

次の図に示すように、クエリとそれに関連付けられた実行プランがあります。



ドロップダウンメニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。
注意: 正しい選択ごとに 1 ポイントが加算されます。

The performance issue stems from the **[answer choice]** operator.

The performance issue can be resolved by adding include columns to the **[answer choice]**.

Microsoft

Select
Index Seek
Key Lookup
Nested Loops

heap
clustered index
nonclustered index

正解:

The performance issue stems from the [answer choice] operator.

Microsoft
jpnshiken.com

Select
Index Seek
Key Lookup
Nested Loops

The performance issue can be resolved by adding include columns to the [answer choice].

heap
clustered index
nonclustered index

Explanation:

Microsoft
jpnshiken.com

The performance issue stems from the [answer choice] operator.

Select
Index Seek
Key Lookup
Nested Loops

The performance issue can be resolved by adding include columns to the [answer choice].

heap
clustered index
nonclustered index

Box 1: Key Lookup

The Key Lookup cost is 99% so that is the performance bottleneck.

Box 2: nonclustered index

The key lookup on the clustered index is used because the nonclustered index does not include the required columns to resolve the query. If you add the required columns to the nonclustered index, the key lookup will not be required.

Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an Azure SQL database.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 98

管理者に通知するためのソリューションを実装する必要があります。このソリューションは監視要件を満たす必要があります。

何をすべきでしょうか？

A. 静的しきい値を持つ Azure Monitor アラート ルールを作成し、そのアラート ルールをアクション グループに割り当てます。

- B.** QueryStoreRuntimeStatistics をログに記録し、Azure イベント ハブにストリームする診断設定を追加します。
- C.** タイムアウトを記録し、Azure イベント ハブにストリームする診断設定を追加します。
- D.** 動的しきい値を持つ Azure Monitor アラート ルールを作成し、そのアラート ルールをアクション グループに割り当てます。

正解: [\(正解を表示します\)](#)

Reference:

<https://azure.microsoft.com/en-gb/blog/announcing-azure-monitor-aiops-alerts-with-dynamic-thresholds/>

質問: 99

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、述べられた目標を達成する可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

Table1という名前のテーブルを含むAzureSynapseAnalytics専用のSQLプールがあります。

container1という名前のAzureData Lake StorageGen2コンテナに取り込まれてロードされるファイルがあります。

ファイルからTable1にデータを挿入し、データを変換することを計画しています。ファイル内のデータの各行は、Table1のサービングレイヤーに1つの行を生成します。

ソースデータファイルがcontainer1にロードされるときに、DateTimeがTable1の追加の列として格納されていることを確認する必要があります。

解決策 :Azure Synapse Analyticsパイプラインでは、ファイルの日時を取得するメタデータの取得アクティビティを使用します。

これは目標を達成していますか？

A. はい

B. いいえ

正解: **B** ([コメントを發表する](#))

Instead use a serverless SQL pool to create an external table with the extra column.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: No is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario asks for: Each question in the series contains a unique solution that might meet the stated goals. That makes No the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 100

複数のデータベースをホストするオンプレミスのMicrosoft SQL Serverインスタンス(SQL1)があります。Azureサブスクリプションには、Azure SQL Managed Instance(SQLM11)が含まれています。Azure Data Studioを使用して、SQLiデータベースをSQLMHにオフライン移行する必要があります。SQL1には何をインストールすればよいのでしょうか？

A. データ移行アシスタント (DMA)

B. セルフホスト型統合ランタイム

C. SQL Server データ ツール (SSDT)

D. Azure Migrate アプライアンス

正解: **B** ([コメントを發表する](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: a self-hosted integration runtime is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario asks for: You need to perform an offline migration of the SQLi databases to SQLMH by using Azure Data Studio. That makes a self-hosted integration runtime the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Data Migration Assistant (DMA) is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why C is Wrong: SQL Server Data Tools (SSDT) is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: an Azure Migrate appliance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

質問: 101

Azure サブスクリプションをお持ちです。

オンプレミスの Microsoft SQL Server インスタンス 10 個を Azure に移行する予定です。

移行後の環境がマルチサーバー管理を使用して管理可能であり、マスター/ターゲット(MSX/TSX)ジョブをサポートしていることを確認する必要があります。ソリューションは管理作業を最小限に抑える必要があります。

マスター サーバー (MSX) とターゲット サーバー (TSX) として選択する必要がある SQL 展開オプションはどれですか？

回答するには、回答エリアで適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area



MSX:


- SQL database
- SQL managed instances
- SQL virtual machines

TSX:

- SQL database
- SQL managed instances
- SQL virtual machines

正解:

Answer Area



Microsoft

MSX:
SQL database
SQL managed instances
SQL virtual machines

TSX:
SQL database
SQL managed instances
SQL virtual machines

Explanation:

Answer Area

MSX:
TSX:

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to ensure that the migrated environment can be managed by using multiserver administration and supports master/target (MSX/TSX) jobs.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 102

Azure サブスクリプションをお持ちです。

Azure Resource Manager (ARM) テンプレートを使用して、Azure SQL マネージド インスタンスをデプロイする必要があります。

ソリューションは次の要件を満たす必要があります。

SQL マネージド インスタンスには一意の ID を割り当てる必要があります。

Azure データセンターの停止が発生した場合でも、SQL マネージド インスタンスが利用可能である必要があります。

テンプレートはどのように完成させるべきでしょうか？ 適切な値を正しいターゲットにドラッグしてください。各値は1回、複数回、または全く使用されない場合があります。コンテンツを表示するには、ペイン間の分割バーをドラッグするか、スクロールする必要がある場合があります。

注意: 正しい選択ごとに 1 ポイントが加算されます。



Explanation:

Values

dnsZonePartner

storageAccountType

SystemAssigned

UserAssigned

zoneRedundant

Answer Area

```

{
  "type": "Microsoft.Sql/managedInstances",
  "identity": {
    "type": "UserAssigned"
  },
  "dependsOn": [
    "[parameters('virtualNetworkName')]"
  ],
  "properties": {
    "administratorLogin": "[parameters('administratorLogin')]",
    "administratorLoginPassword": "[parameters('administratorLoginPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets',
    parameters('virtualNetworkName'), parameters('subnetName'))]",
    "storageSizeInGB": 8192,
    "vCores": 80, "licenseType": "BasePrice",
    "storageAccountType": "true"
  }
}

```



Microsoft

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to deploy an Azure SQL managed instance by using an Azure Resource Manager (ARM) template.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 103

次の表に示すリソースを含む Azure サブスクリプションがあります。

Name	Type
App1	Azure web app
db1	Azure SQL database in the serverless tier

App1 は、長時間操作が行われなかった後に db1 にアクセスしようとする、一時的な接続エラーとタイムアウトが発生します。App1 で発生している問題を解決するために、当面のコストを考慮せずに、できるだけ早く db1 を修正する必要があります。どうすればよいでしょうか？

- A. db1 に割り当てられている vCore の数を増やします。
- B. db1 の自動一時停止遅延を無効にします。
- C. db1 の自動一時停止遅延を減らします。
- D. db1 の自動チューニングを有効にします。

正解: ([正解を表示します](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: Enable automatic tuning for db1. matches the expected DP-300 administration action.

Automatic tuning uses built-in recommendations and verification to create/drop indexes or force last known good plans where supported. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to modify db1 to resolve the issues experienced by App1 as soon as possible, without considering immediate costs.

Why A is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It does not expose the required metric, query history, wait

/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Disable auto-pause delay for db1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Decrease the auto-pause delay for db1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 104

SQLMI1 という名前の Azure SQL マネージド インスタンスと Workspace1 という名前の Log Analytics ワークスペースを含む Azure サブスクリプションがあります。

SQLMI1 のパフォーマンス メトリックを収集し、そのメトリックを Workspace1 にストリーミングする必要があります。

A. SQLMI1 にプライベート エンドポイント接続を作成します。

B. Workspace1 を使用するように Azure SQL Analytics を構成します。

C. SQLMI1 のコンピューター + ストレージ設定を変更します。

D. SQLMI1 の診断設定を変更します。

正解: ([正解を表示します](#))

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why D is Correct: Modify the diagnostic settings for SQLMI1. matches the expected DP-300 administration action. Modify the diagnostic settings for SQLMI1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to collect performance metrics for SQLMI1 and stream the metrics to Workspace1.

Why A is Wrong: A private endpoint exposes an Azure service through a private IP address in a virtual network, avoiding public endpoint dependency. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why B is Wrong: Configure Azure SQL Analytics to use Workspace1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Modify the Computer + storage settings for SQLMI1. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 105

次のAzureDataFactoryパイ プラインがあります。

System1からデータを取り込む

System2からデータを取り込む

寸法を入力します

事実を入力する

System1からのデータの取り込みとSystem2からのデータの取り込みには依存関係がありません。ディメンションの取り込みは、System1からデータを取り込み、System2からデータを取り込む後に実行する必要があります。Populate Factsは、PopulateDimensionsパイ プラインの後に実行する必要があります。すべてのパイ プラインは8時間ごとに実行する必要があります。

パイプラインの実行をスケジュールするにはどうすればよいですか？

- A. 4つのパイプラインすべてにスケジュールトリガーを追加します。
- B. 4つのパイプラインすべてにイベントトリガーを追加します。
- C. 4つのパイプラインを含む親パイプラインを作成し、イベントトリガーを使用します。
- D. 4つのパイプラインを含む親パイプラインを作成し、スケジュールトリガーを使用します。

正解: ([正解を表示します](#))

Reference:

<https://www.mssqltips.com/sqlservertip/6137/azure-data-factory-control-flow-activities-overview/> Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why D is Correct: Create a parent pipeline that contains the four pipelines and use a schedule trigger. matches the expected DP-300 administration action. Create a parent pipeline that contains the four pipelines and use a schedule trigger. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Populate Dimensions must execute after Ingest Data from System1 and Ingest Data from System2.

Why A is Wrong: Add a schedule trigger to all four pipelines. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why B is Wrong: Add an event trigger to all four pipelines. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: Create a parent pipeline that contains the four pipelines and use an event trigger. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

質問: 106

リソース ガバナーが有効になっており、App1 と App2 という 2 つのアプリによって使用されている SQLMI1 という名前の Azure SQL マネージド インスタンスがあります。

App1 に割り当てることができる CPU およびメモリ リソースを制限するように SQLMI1 を構成する必要があります。

どの 4 つのアクションを順番に実行する必要がありますか? 回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions	Answer Area
Create a workload group.	⏪ ⏩ ⏴ ⏵ Microsoft
Create a user-defined classifier function.	
Modify Resource Governor.	
Create a contained database user	
Create a resource pool.	

正解:

Actions

- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.
- Create a contained database user.
- Create a resource pool.

Answer Area

- Create a resource pool.
- Create a workload group.
- Create a user-defined classifier function.
- Modify Resource Governor.

Explanation:

Create a resource pool.

Create a workload group.

Create a user-defined classifier function.

Modify Resource Governor.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/resource-governor?view=sql-server-ver15>

<https://docs.microsoft.com/en-us/sql/relational-databases/resource-governor/create-and-test-a-classifier-user-defined-function?view=sql-server-ver15> Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to configure SQLMI1 to limit the CPU and memory resources that can be allocated to App1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> 503問、30%ディスカウント、特別な割引コード: **JPNshiken**」

質問: 107

タスク 7

Azure AD ID から db1 にデータベース ユーザーを作成する自動化 Runbook を作成する予定です。
新しいデータベース ユーザーの作成をサポートするには、sq1370O6895 を構成する必要があります。

正解:
See the explanation part for the complete Solution.
Explanation:

To configure sq1370O6895 to support the creation of new database users from Azure AD identities, you need to do the following steps:

- * Set up a Microsoft Entra tenant and associate it with your Azure subscription. You can use the Microsoft Entra portal or the Azure portal to create and manage your Microsoft Entra users and groups12.
- * Configure a Microsoft Entra admin for sq1370O6895. You can use the Azure portal or the Azure CLI to set a Microsoft Entra user as the admin for the server34. The Microsoft Entra admin can create other database users from Microsoft Entra identities5.
- * Connect to db1 using the Microsoft Entra admin account and run the following Transact-SQL statement to create a new database user from a Microsoft Entra identity: CREATE USER [Microsoft Entra user name] FROM EXTERNAL PROVIDER;6 You can replace the Microsoft Entra user name with the name of the user or group that you want to create in the database.
- * Grant the appropriate permissions to the new database user by adding them to a database role or granting them specific privileges. For example, you can run the following Transact-SQL statement to add the new user to the db_datareader role: ALTER ROLE db_datareader ADD MEMBER [Microsoft Entra user name]; These are the steps to configure sq1370O6895 to support the creation of new database users from Azure AD identities.

質問: 108

factSalesという名前のテーブルを含むAzureSQLデータベースがあります。FactSalesには、次の表に示す列が含まれています。

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

FactSalesには60億行あり、バッチプロセスを使用して毎晩読み込まれます。
どのタイプの圧縮がデータベースに最大のスペース削減を提供しますか？

- A. ページ圧縮
- B. 行の圧縮
- C. 列ストアの圧縮
- D. 列ストアアーカイブの圧縮

正解: (正解を表示します)
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.
Reference:
<https://docs.microsoft.com/en-us/sql/relational-databases/data-compression/data-compression>

質問: 109

Azure サブスクリプションをお持ちです。
Azure Virtual Machines 上に SQL Server インスタンスを展開する必要があります。ソリューションは以下の要件を満たす必要があります。
* ICPS、容量、スループットなどのカスタム パフォーマンス構成をサポートする必要があります。
* コストを最小限に抑える必要がある

ソリューションに含めるべきディスクの種類はどれですか？

- A. プレミアム SSD v2
- B. プレミアムSSD
- C. ウルトラSSD
- D. 標準SSD

正解: **A** ([コメントを发表する](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: Premium SSD v2 is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. In this scenario, the important constraint is: You need to deploy an Instance of SQL Server on Azure Virtual Machines. Premium SSD v2 satisfies that constraint without adding an unrelated service or manual process.

Why B is Wrong: Premium SSD is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Ultra SSD is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: Standard SSD is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: **110**

DB1という名前のオンプレミスのMicrosoft SQL Server 2022データベースがあります。

あなたは、以下のTransact-SQLスクリプトに基づいてデータベース監査仕様を作成する予定です。

```
CREATE SERVER AUDIT Security_Audit TO FILE (FILEPATH = 'C:\Program Files\Microsoft SQL
Server\MSSQL13.MSSQLSERVER\MSSQL\DATA');
GO

USE DB1;
GO

CREATE DATABASE AUDIT SPECIFICATION Audit_Card_Tables
FOR SERVER AUDIT Security_Audit ADD (
    SELECT, INSERT ON Payment.Creditcard BY dbo
)
WITH (STATE = ON);
GO
```



以下の各記述について、正しい場合は「はい」を、そうでない場合は「いいえ」を選択してください。

注：正解ごとに1ポイントが加算されます。

Answer Area

Microsoft

Statements	Yes	No
The server audit will be enabled.	☐	☐
The database audit will audit all SELECT statements to the Payment.Creditcard table.	☐	☐
You must be granted at least the CONTROL DATABASE permission to perform the task.	☐	☐

正解:

Answer Area

Microsoft

Statements	Yes	No
The server audit will be enabled.	☐	☒
The database audit will audit all SELECT statements to the Payment.Creditcard table.	☒	☐
You must be granted at least the CONTROL DATABASE permission to perform the task.	☒	☐

Explanation:

Answer Area

Microsoft

Statements	Yes	No
The server audit will be enabled.	☐	☒
The database audit will audit all SELECT statements to the Payment.Creditcard table.	☒	☐
You must be granted at least the CONTROL DATABASE permission to perform the task.	☒	☐

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have an on-premises Microsoft SQL Server 2022 database named DB1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

マージ レプリケーションを使用する SQL1 という名前のオンプレミスの Microsoft SQL Server 2019 データベースがあります。SQL1 を Azure に移行する必要があります。どのサービスを使用する必要がありますか？

- A. Azure SQL エッジ
- B. Azure SQL データベース
- C. Azure 仮想マシン上の SQL Server
- D. Azure SQL マネージド インスタンス

正解: **C** ([コメントを發表する](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why C is Correct: SQL Server on Azure Virtual Machines matches the expected DP-300 administration action. SQL Server on Azure Virtual Machines is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to migrate SQL1 to Azure.

Why A is Wrong: Azure SQL Edge is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why B is Wrong: Azure SQL Database is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: Azure SQL Managed instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

質問: 112

Web層、アプリケーション層、およびMicrosoft SQLServer層を含むApp1という名前のオンプレミス多層アプリケーションがあります。すべての層はHyper-V仮想マシンで実行されます。

新しいディザスタリカバリ計画では、ビジネスに不可欠なすべてのアプリケーションをAzureにリカバリできる必要があります。

App1のデータベース層をAzureにフェイルオーバーするソリューションを推奨する必要があります。このソリューションは、現在の環境に影響を与えることなく、Azureへのフェイルオーバーをテストする機能を提供する必要があります。

推奨事項には何を含める必要がありますか？

- A. Azureバックアップ
- B. Azure Information Protection
- C. WindowsServerフェールオーバークラスター
- D. Azureサイトの回復

正解: **D** ([コメントを發表する](#))

Reference:

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-test-failover-to-azure> Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: Azure Site Recovery is correct because it is the feature whose normal purpose matches the stated requirement. Azure Site Recovery is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: Your new disaster recovery plan requires that all business-critical applications can be recovered to Azure.

Why A is Wrong: Azure Backup is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why B is Wrong: Azure Information Protection is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: Windows Server Failover Cluster is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario 's failure model.

質問: 113

DB1 という名前の Azure SCX データベースを保存します。

D61 内のテーブルのデータとインデックスの断片化情報を照会する必要があります。

どのコマンドを実行する必要がありますか？

A. `sys.dm_db_index_physical_stats`

B. `DBCC SHOWCONTIG`

C. `DBCC CHECKALLOC`

D. `sys.dm_db_index_usage_stats`

正解: **A** ([コメントを发表する](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: is correct because it is the feature whose normal purpose matches the stated requirement. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: You need to query the fragmentation information of data and indexes for the tables in D61.

Why B is Wrong: is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: 114

PaaS プロトタイプに基づいて、どの Azure SQL Database コンピューティング レベルを使用する必要がありますか？

A. ビジネスクリティカル 4 仮想コア

B. ハイパースケール

C. 汎用 v-vCore

D. サーバーレス

正解: ([正解を表示します](#))

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

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why D is Correct: The serverless tier automatically scales compute and can pause during inactive periods, making it suitable for intermittent workloads with bursts and idle windows. The scenario asks for: Based on the PaaS prototype, which Azure SQL Database compute tier should you use? That makes Serverless the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. That design keeps capacity available for a different kind of workload and does not match the cost behavior expected from an idle-and-burst database pattern.

Why B is Wrong: Hyperscale is designed for very large databases and fast scale-out storage architecture, not merely for intermittent CPU bursts. That design keeps capacity available for a different kind of workload and does not match the cost behavior expected from an idle-and-burst database pattern.

Why C is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It applies to a different Azure data-platform design than the one the question is testing.

Verification Note: The original listed answer was A; the corrected answer is D because The prototype shows intermittent spikes with idle periods; the serverless compute tier is designed for intermittent usage and automatic pause/resume cost savings.

質問: 115

次のTransact-SQLクエリがあります。

```
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;
```

クエリによって返されるどの列が各ファイルの空き領域を表しますか？

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

正解: ([正解を表示します](#))

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/> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: ColumnC is correct because it is the feature whose normal purpose matches the stated requirement. ColumnC is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: Which column returned by the query represents the free space in each file?

Why A is Wrong: ColumnA is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: ColumnB is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: ColumnD is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 116

汎用サービス層に、DB1という名前のAzure SQLデータベースがあります。

DB1のパフォーマンス指標を以下の図に示します。



Log IOの割合を減らす必要があります。解決策はコストを最小限に抑えるものでなければなりません。

あなたはすべきですか？

- A. vCoreの数を増やします。
- B. B. リカバリーモデルをシンプルに変更します。
- C. チェックポイント操作を実行します。

D. サービス階層をビジネスクリティカルに変更します。

正解: ([正解を表示します](#))

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why D is Correct: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads.

This fits the case because the requested outcome is: You need to reduce the Log 10 percentage. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: The vCore purchasing model separates compute and storage sizing and supports Azure Hybrid Benefit for eligible SQL Server licenses. It applies to a different Azure data-platform design than the one the question is testing.

Why B is Wrong: B. Change Recovery model to Simple. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Perform a checkpoint operation. is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

質問: **117**

Azure サブスクリプションをお持ちです。

Azure Marketplace イメージを使用して Azure Virtual Machines 上に SQL Server インスタンスを展開する予定の場合、SQL Server IaaS Agent 拡張機能 (SqlIaaSExtension) を登録する必要があります。ソリューションは以下の要件を満たす必要があります。

* SQL Server の重要な更新プログラムを自動的にインストールします。

* 仮想マシンへのパフォーマンスの影響を最小限に抑えます。

どの管理モードを選択する必要がありますか？

A. フル

B. 軽量

C. エージェントなし

正解: ([正解を表示します](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why A is Correct: full matches the expected DP-300 administration action. full is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You plan to deploy an instance of SQL Server on Azure Virtual Machines by using an Azure Marketplace image. You need to register the SQL Server IaaS Agent extension (SqlIaaSExtension).

Why B is Wrong: lightweight is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why C is Wrong: NoAgent is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: **118**

Azure SQL マネージド インスタンスがあります。

SQL エージェント ジョブの電子メール通知を有効にする必要があります。

何をすべきでしょうか？

A. エージェント XP オプションを使用します。

B. SQL Server エージェントを有効にします。

C. sp_configure コマンドを実行します。

D. sp set_agent properties コマンドを実行します。

正解: C ([コメントを发表する](#))

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why C is Correct: Run the sp_configure command. can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. This fits the case because the requested outcome is: You need to enable SQL Agent Job email notifications. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Use the Agent XPs option. can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why B is Wrong: SQL Server Agent schedules and runs SQL Server jobs; in Azure SQL Database, Elastic Jobs or Azure Automation are often used instead. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

Why D is Wrong: Run the sp_set_agent_properties command. can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

質問: 119

Azure Synapse Analytics サーバーレス SQL プールを使用して、Azure Data Lake Storage Gen2 アカウント内のバス運賃データの探索的分析を実行しています。

次の図に示す Transact-SQL クエリを実行します。

```
SELECT
    payment_type,
    SUM(fare_amount) AS fare_total
FROM OPENROWSET(
    BULK 'csv/busfare/tripdata_2020*.csv',
    DATA_SOURCE = 'BusData',
    FORMAT = 'CSV', PARSER_VERSION = '2.0',
    FIRSTROW = 2
)
WITH (
    payment_type INT 10,
    fare_amount FLOAT 11
) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。

The query results include only **[answer choice]** in the csv/busfare folder.

The query assumes that the first row in a CSV file is **[answer choice]** row.

Microsoft ▼ CSV files in the tripdata_2020 subfolder files that have files names beginning with "tripdata_2020" CSV files that have file names containing "tripdata_202" CSV files that have file named beginning with "tripdata_2020"

▼ a header a data an empty
--

正解:

The query results include only **[answer choice]** in the csv/busfare folder.

The query assumes that the first row in a CSV file is **[answer choice]** row.

Microsoft ▼ CSV files in the tripdata_2020 subfolder files that have files names beginning with "tripdata_2020" CSV files that have file names containing "tripdata_202" CSV files that have file named beginning with "tripdata_2020"

▼ a header a data an empty
--

Explanation:

The query results include only **[answer choice]** in the csv/busfare folder.

The query assumes that the first row in a CSV file is **[answer choice]** row.

Microsoft ▼ CSV files in the tripdata_2020 subfolder files that have files names beginning with "tripdata_2020" CSV files that have file names containing "tripdata_202" CSV files that have file named beginning with "tripdata_2020"

▼ a header a data an empty
--

Box 1: CSV files that have file named beginning with " tripdata_2020 "

Box 2: a header

FIRSTROW = ' first_row '

Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset> Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 120

Vault1 という名前の Azure キー コンテナーと、Azure Virtual Machines インスタンス上の SQL Server (SQL1) があり、SQL 1 は DB1 という名前のデータベースをホストしています。

Vault1 のキーを使用するには、D61 で透過的なデータ暗号化 (TDE) を構成する必要があります。どの 4 つのアクションを順番に実行する必要がありますか? 回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序で並べます。

Actions

- On SQL1, create a symmetric key.
- On SQL1, create a login from the symmetric key.
- Create an Azure AD service principal and grant the service principal permissions for Vault1.
- On SQL1, create an asymmetric key.
- On SQL1, create a cryptographic provider and a Microsoft SQL Server credential
- On SQL1, create a login from the asymmetric key.

Answer Area

正解:

Actions

- On SQL1, create a symmetric key.
- On SQL1, create a login from the symmetric key.
- Create an Azure AD service principal and grant the service principal permissions for Vault1.
- On SQL1, create an asymmetric key.
- On SQL1, create a cryptographic provider and a Microsoft SQL Server credential
- On SQL1, create a login from the asymmetric key.

Answer Area

- Create an Azure AD service principal and grant the service principal permissions for Vault1.
- On SQL1, create an asymmetric key.
- On SQL1, create a cryptographic provider and a Microsoft SQL Server credential
- On SQL1, create a login from the asymmetric key.

Explanation:

Actions

- On SQL1, create a symmetric key.
- On SQL1, create a login from the symmetric key.

Answer Area

- Create an Azure AD service principal and grant the service principal permissions for Vault1.
- On SQL1, create an asymmetric key.
- On SQL1, create a cryptographic provider and a Microsoft SQL Server credential.
- On SQL1, create a login from the asymmetric key.

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to configure Transparent Data Encryption (TDE) on D61 to use a key in Vault1 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.

Why the alternate choices are Wrong: The alternate sequence fails because key custody, application identity, and column encryption must be prepared in the correct order before the wizard can protect sensitive columns.

質問: 121

AzureSQLデータベースがあります。データベースには、列ストアインデックスを使用し、アクセス頻度が低いテーブルが含まれています。

列ストアのアーカイブ圧縮を有効にします。

構成の2つの可能な結果は何ですか？それぞれの正解は完全な解決策を提示します。

注：正しい選択はそれぞれ1ポイントの価値があります。

A. インデックスを使用するクエリは、より多くのディスクI/Oを消費します。

B. インデックスを使用するクエリでは、取得するデータページが少なくなります。

C. インデックスはより多くのディスクスペースを消費します。

D. インデックスはより多くのメモリを消費します。

E. インデックスを使用するクエリは、より多くのCPUリソースを消費します。

正解: B,E ([コメントを發表する](#))

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.

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why B and E is Correct: Queries that use the index will retrieve fewer data pages. matches the expected DP-

300 administration action. Queries that use the index will retrieve fewer data pages. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: Each correct answer presents a complete solution. CPU percentage is important for database compute monitoring, but it is not the Stream Analytics sizing metric when the question is about streaming units. The scenario asks for: Each correct answer presents a complete solution. That makes Queries that use the index will consume more CPU resources. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Queries that use the index will consume more disk I/O. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.


Why C is Wrong: The index will consume more disk space. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: The index will consume more memory. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 122

index1 という名前の非クラスター化インデックスを含む DBI という名前の Azure SQL データベースがあります。
エンドユーザーから、index1 を使用するとクエリが遅くなるという報告があります。
インデックスに対して実行されている操作を識別する必要があります。
どの動的管理ビューを使用する必要がありますか？



- A. オプションA
- B. オプションB
- C. オプションC
- D. オプションD

正解: ([正解を表示します](#))

Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why D is Correct: Option D is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation.

This fits the case because the requested outcome is: You need to identify the operations that are being performed on the index.

The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Option A is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Option B is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Option C is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 123

AzureSQLデータベースがあります。

ユーザーは、ストアードプロシージャの実行が通常より遅いと報告しています。リグレッションされたクエリがパフォーマンスの問題を引き起こしていると思われます。

クエリ実行プランを表示して、リグレッションされたクエリが問題の原因であるかどうかを確認する必要があります。ソリューションは労力を最小限に抑える必要があります。

何を使うべきですか？

- 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

正解: **C** ([コメントを发表する](#))

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>

質問: 124

あなたはAzureサブスクリプションをお持ちです。

BACPACファイルを使用してAzure SQLデータベースをデプロイする必要があります。どのコマンドを実行すればよいですか？

A. az sql db インポート

B. az sql db copy

C. az sql db op

D. az sql db create

正解: ([正解を表示します](#))

Basic Concept: This question tests configure and manage automation of tasks in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why A is Correct: az sql db import can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule.

The scenario asks for: You need to deploy an Azure SQL database by using a BACPAC file. That makes az sql db import the option that satisfies the required Azure SQL layer and operational outcome.

Why B is Wrong: az sql db copy can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why C is Wrong: az sql db op can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It does not provide the required repeatable execution, scheduling, or deployment automation requirement, so it would not complete the administrative workflow reliably.

Why D is Wrong: az sql db create can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

質問: 125

Windows Server 2019 を実行する VM1 と VM2 という 2 つの Azure 仮想マシンがあります。VM1 と VM2 はそれぞれ、デフォルトの Microsoft SQL Server 2019 インスタンスをホストしています。VM1 には DB1 というデータベースが含まれており、これは D:\DB1.bak というファイルにバックアップされています。

次の構成を持つ Always On 可用性グループを展開する予定です。

* VM1 は DB1 のプライマリ レプリカをホストします。

* VM2 は DB1 のセカンダリ レプリカをホストします。

可用性グループ用に VM2 上にセカンダリ データベースを準備する必要があります。

Transact-SQL ステートメントをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

▼ DATABASE MyDB1

BACKUP
CREATE
RESTORE

FROM DISK = 'D:\DB1.bak'

WITH

GO

▼
NORECOVERY
RECOVERY
STANDBY

正解:

▼ DATABASE MyDB1

BACKUP
CREATE
RESTORE

FROM DISK = 'D:\DB1.bak'

WITH

GO

▼
NORECOVERY
RECOVERY
STANDBY

Explanation:



Reference:

<https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/manually-prepare-a-secondary-database-for-an-availability-group-sql-server?view=sql-server-ver15> Basic Concept:

This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to prepare the secondary database on VM2 for the availability group.

Why the alternate choices are Wrong: The alternate choices do not satisfy the stated availability design because quorum and listener resources must match Azure failover-cluster behavior and the required SLA.

質問: 126

Azure Virtual Machines 上に SQL1 という名前の SQL Server インスタンスがあります。

SQL1 には、Microsoft SQL Server イベントをキャプチャする session! という名前の拡張イベント セッションが含まれています。

セッション イベントを、Event Tracing for Windows (ETW) によってキャプチャされたイベントと関連付ける必要があります。

セッション1では何をすべきでしょうか？

- A. セッション イベント フィルターの設定を変更します。
- B. ターゲットを追加します。
- C. アクションを追加します。
- D. セッション データ ストレージの指定設定を変更します。

正解: ([正解を表示します](#))

Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: Add a target. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to correlate the session events with events captured by Event Tracing for Windows (ETW). The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: A session window groups events separated by periods of inactivity, which is useful for user sessions rather than fixed continuous aggregation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Add an action. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: A session window groups events separated by periods of inactivity, which is useful for user sessions rather than fixed continuous aggregation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 127

DB1という名前のAzureSQLデータベースがあります。DB1に接続しているときにクエリを実行します。クエリの実際の実行プランを確認し、クエリによって参照されるテーブルにインデックスを追加します。クエリの以前の実際の実行プランをライブクエリ統計と比較する必要があります。Microsoft SQL Server Management Studio (SSMS)で最初に何をする必要がありますか？

- A. DB1の場合、クエリストアのQUERY_CAPTURE_MODEをAllに設定します。
- B. SET SHOWPLAN_ALL Transact-SQLステートメントを実行します。
- C. 実際の実行プランを保存します。
- D. DB1のクエリストアを有効にします。

正解: ([正解を表示します](#))

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.

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why C is Correct: Save the actual execution plan. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario asks for: You need to compare the previous actual execution plan for the query to the Live Query Statistics. That makes Save the actual execution plan. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why B is Wrong: Run the SET SHOWPLAN_ALL Transact-SQL statement. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why D is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: 128

あなたはAzureサブスクリプションをお持ちです。

サブスクリプション内にDB1という名前の新しいデータベースを作成する必要があります。ソリューションは以下の要件を満たす必要があります。

DB1で作成されたテーブルは、データが変更された際に自動的に暗号化による証明を提供するようにしてください。

監査担当者が、データが改ざんされていないことを検証できるようにしてください。

どのT-SQLコマンドを実行すればよいですか？

- A. サーバー証明書によるAES_256暗号化アルゴリズムを使用してデータベース暗号化キーを作成します

B. CREATE DATABASE DB1 WITH TRUSTWORTHY

C. CREATE DATABASE DB1 WITH LEDGER = ON

D. CREATE DATABASE DB1 WITH EDITION = ' GeneralPurpose ', SERVICE_OBJECTIVE = ' GP_Gen5_2 '

正解: ([正解を表示します](#))

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why C is Correct: CREATE DATABASE DB1 WITH LEDGER = ON is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. In this scenario, the important constraint is: You need to create a new database named DB1 in the subscription. CREATE DATABASE DB1 WITH LEDGER = ON satisfies that constraint without adding an unrelated service or manual process.

Why A is Wrong: Certificates support encryption and authentication, but the correct certificate-related object depends on whether the task is TLS, Always Encrypted, or user authentication. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why B is Wrong: CREATE DATABASE DB1 WITH TRUSTWORTHY is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: CREATE DATABASE DB1 WITH EDITION = ' GeneralPurpose ', SERVICE_OBJECTIVE = ' GP_Gen5_2 ' is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

質問: 129

Windows Server 2022 を実行する Server1 という名前の Azure 仮想マシンがあります。Server1 には、SQL1 という名前の Microsoft SQL Server 2022 のインスタンスと、DB1 という名前のデータベースが含まれています。

SQL1 のマスター データベースにマスター キーを作成します。

DB1 の暗号化されたバックアップを作成する必要があります。

何をすべきでしょうか？

A. DB1 に対称キーを作成します。

B. Server1 で視覚化ベースのセキュリティ (VBS) を有効にします。

C. DB1 に証明書を作成します。

D. SQL1 のマスター データベースに証明書を作成します。

正解: **D** ([コメントを发表する](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: Create a certificate in the master database of SQL1. is correct because it is the feature whose normal purpose matches the stated requirement. Certificates support encryption and authentication, but the correct certificate-related object depends on whether the task is TLS, Always Encrypted, or user authentication. The scenario wording points to that specific behavior: You need to create an encrypted backup of DB1.

Why A is Wrong: Create a symmetric key in DB1. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why B is Wrong: Enable visualization-based security (VBS) on Server1. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why C is Wrong: Certificates support encryption and authentication, but the correct certificate-related object depends on whether the task is TLS, Always Encrypted, or user authentication. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to create an encrypted backup of DB1.

質問: 130

DB1という名前のAzureSQLデータベースがあります。

Azureポータルのクエリエディターを使用して、クエリの推定実行プランを表示する必要があります。

あなたは最初に何をすべきですか？

- A. set showplan_all Transact-SQLステートメントを実行します。
- B. DB1の場合、クエリストアのQUERY_CAPTURE_MODEをAllに設定します。
- C. set forceplan Transact-SQLステートメントを実行します。
- D. DB1のクエリストアを有効にします。

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/set-showplan-all-transact-sql?view=sql-server-ver15> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why A is Correct: Run the set showplan_all Transact-SQL statement. matches the expected DP-300 administration action. Run the set showplan_all Transact-SQL statement. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to display the estimated execution plan of a query by using the query editor in the Azure portal.

Why B is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why C is Wrong: Run the set forceplan Transact-SQL statement. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

Why D is Wrong: Query Store captures query text, plans, and runtime statistics, making it central for reviewing query history and troubleshooting plan regressions. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.


質問: 131

DB1 という名前のデータベースが含まれる、Server1 という名前のオンプレミスの Microsoft SQL Server 2016 サーバーがあります。

Azure Database Migration Service を使用して、DB1 から Azure SQL Database マネージド インスタンスへのオンライン移行を実行する必要があります。

DB1 のバックアップをどのように構成すればよいですか？ 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD

Explanation:

Backup type:

Full and log backups only
Full backup only
Log backup only

Backup option:

WITH CHECKSUM
WITH NOINIT
WITH UNLOAD

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.

Reference:

<https://docs.microsoft.com/en-us/azure/dms/known-issues-azure-sql-db-managed-instance-online> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to perform an online migration of DB1 to an Azure SQL Database managed instance by using Azure Database Migration Service.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 132

Azure SQL データベースがあります。


遅延されたクエリ実行が RESOURCE 待機に関連付けられているかどうかを識別する必要があります。

Transact-SQL ステートメントをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。



正解:



Explanation:



Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to identify whether a delayed query execution is associated to a RESOURCE wait.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 133

Azure SQL マネージド インスタンスがあります。

クエリ プランの最後の実行とそのランタイム統計を収集する必要があります。ソリューションは、現在実行中のクエリへの影響を最小限に抑える必要があります。

あなたは何をするべきか？

- A. 推定実行計画を生成します。
- B. 実際の実行計画を生成します。
- C. sys.dm_exec_query_plan_stats を実行します。
- D. ライブ クエリ統計を生成します。

正解: C ([コメントを發表する](#))

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-exec-query-plan-stats-transact-sql?view=sql-server-ver15> Basic Concept: This question tests performance optimization for SQL workloads, including Query Store, automatic tuning, index/statistics maintenance, and engine-level resource behavior.

Why C is Correct: Run sys.dm_exec_query_plan_stats. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. This fits the case because the requested outcome is: You need to gather the last execution of a query plan and its runtime statistics. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Generate an estimated execution plan. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why B is Wrong: Generate an actual execution plan. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why D is Wrong: Generate Live Query Statistics. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: 134

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、述べられた目標を達成する可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

Azure仮想マシン上にSQLServer2019があります。

SQLServerインスタンスでのクエリのパフォーマンスの問題のトラブルシューティングを行っています。

詳細情報を収集するには、sys.dm_exec_requestsにクエリを実行し、待機タイプがPAGELATCH_UPであり、wait_resourceが2 :3 :905856であることを確認します。

システムパフォーマンスを改善する必要があります。

解決策 :マスターデータベースのデータ ファイルを10%自動拡張するように変更します。

これは目標を達成していますか？

- A. はい
- B. いいえ

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention> Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why B is Correct: No is correct because it is the feature whose normal purpose matches the stated requirement. No is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: Each question in the series contains a unique solution that might meet the stated goals.

Why A is Wrong: Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 135

顧客がデータベースオブジェクトを作成できるようにするためのソリューションを推奨する必要があります。そのソリューションはビジネス目標を満たすものでなければなりません。

推薦書には何を含めるべきでしょうか？

- A. 顧客ごとに、既存のスキーマに顧客の ddl_admin を付与します。
- B. 顧客ごとに追加のスキーマを作成し、顧客に新しいスキーマへの ddl_admin を付与します。
- C. 顧客ごとに追加のスキーマを作成し、顧客に新しいスキーマへの db_writer を付与します。
- D. 顧客ごとに、既存のスキーマに顧客の db_writer を付与します。

正解: B ([コメントを发表する](#))

Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: For each customer, create an additional schema and grant the customer ddl_admin to the new schema. matches the expected DP-300 administration action. For each customer, create an additional schema and grant the customer ddl_admin to the new schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement.

The question is not asking for a general Azure capability; it is asking for the feature that produces this result:

You need to recommend a solution to ensure that the customers can create the database objects.

Why A is Wrong: For each customer, grant the customer ddl_admin to the existing schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why C is Wrong: For each customer, create an additional schema and grant the customer db_writerto the new schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross- region continuity target stated in the question, even if it is valid in a different availability design.

Why D is Wrong: For each customer, grant the customer db_writerto the existing schema. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

質問: 136

10個のパイプラインを含むAzureDataFactoryがあります。

各パイプラインに、取り込み、変換、またはロードのいずれかの主な目的でラベルを付ける必要があります。Data Factoryの監視エクスペリエンスを使用する場合、ラベルはグループ化とフィルタリングに使用できる必要があります。

各パイプラインに何を追加する必要がありますか？

- A. 注釈
- B. リソースタグ
- C. 実行グループID
- D. ユーザープロパティ
- E. 相関ID

正解: A ([コメントを发表する](#))

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag. You can define tags so you can see their performance or find errors faster.

Reference:
<https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/>
Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.
Why A is Correct: an annotation matches the expected DP-300 administration action. an annotation is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to label each pipeline with its main purpose of either ingest, transform, or load.
Why B is Wrong: a resource tag is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.
Why C is Wrong: a run group ID is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.
Why D is Wrong: a user property is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.
Why E is Wrong: a correlation ID is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

有効的な**DP-300J**問題集はJPNTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTest.comは今最新**DP-300J**試験問題集を提供します。JPNTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: **137**
Pool1という名前のAzureSynapse Analytics専用SQLプールと、DB1という名前のデータベースがあります。DB1には、Tableという名前のファクトテーブルが含まれています。
表1でデータスキューの範囲を特定する必要があります。
Synapse Studioで何をすべきですか？
A. Pool1に接続し、sys.dm_pdw_nodes_db_partition_statsを照会します。
B. 組み込みプールに接続し、DBCCCHECKALLOCを実行します。
C. Pool1に接続し、DBCCCHECKALLOCを実行します。
D. 組み込みプールに接続し、sys.dm_pdw_nodes_db_partition_statsにクエリを実行します。
正解: ([正解を表示します](#))
Use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data.
Reference:
<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.
Why D is Correct: The pdw_nodes dynamic management views apply to Azure Synapse Analytics or Parallel Data Warehouse nodes, not ordinary Azure SQL Database troubleshooting. This fits the case because the requested outcome is: You need to identify the extent of the data skew in Table1. The selected feature addresses that outcome directly rather than relying on a workaround.
Why A is Wrong: The pdw_nodes dynamic management views apply to Azure Synapse Analytics or Parallel Data Warehouse nodes, not ordinary Azure SQL Database troubleshooting. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.
Why B is Wrong: Connect to the built-in pool and run DBCC CHECKALLOC. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait

/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why C is Wrong: Connect to Pool1 and run DBCC CHECKALLOC. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

質問: 138

Pool1 という名前の Azure Synapse Analytics 専用 SQL プールと、Account1 という名前の Azure Data Lake Storage Gen2 アカウントがあります。

外部テーブルを使用して Account1 内のファイルにアクセスする予定です。

外部テーブルを作成するときに参照できるデータ ソースを Pool1 に作成する必要があります。

Transact-SQL ステートメントをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

```
CREATE EXTERNAL DATA SOURCE source1
```

```
WITH
```

```
( LOCATION = 'https://account1. .core.windows.net'
```

blob
dfs
table

PUSHDOWN = ON
TYPE = BLOB_STORAGE
TYPE = HADOOP

正解:

```
CREATE EXTERNAL DATA SOURCE source1
```


```
WITH
```

```
( LOCATION = 'https://account1. .core.windows.net',
```

blob
dfs
table

PUSHDOWN = ON
TYPE = BLOB_STORAGE
TYPE = HADOOP

Explanation:



Box 1: blob

The following example creates an external data source for Azure Data Lake Gen2 CREATE EXTERNAL DATA SOURCE YellowTaxi WITH (LOCATION = 'https://azureopendatastorage.blob.core.windows.net/nyctlc/yellow/ ', TYPE = HADOOP) Box 2: HADOOP Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to create a data source in Pool1 that you can reference when you create the external table.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 139

db1 という名前のデータベースをホストする、SQL1 という名前のオンプレミスの Microsoft SQL Server 2019 インスタンスがあります。

MI1 という名前の Azure SQL マネージド インスタンスと storage1 という名前の Azure ストレージ アカウントを含む Azure サブスクリプションがあります。

db1 を storage1 にバックアップできることを確認する必要があります。ソリューションは以下の要件を満たす必要があります。

* ブロック BLOB ストレージを使用します。

* セキュリティを最大限に高めます。

A. 共有アクセス署名 (SAS) を生成する

B. インフラストラクチャ暗号化を有効にします。

C. アクセス ポリシーを作成します。

D. ストレージキーをローテーションする

正解: (正解を表示します)

Basic Concept: This question tests secure database administration, where the control must match the data state, access boundary, identity model, or compliance requirement.

Why B is Correct: Enable infrastructure encryption. is correct because it is the feature whose normal purpose matches the stated requirement. Enable infrastructure encryption. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario wording points to that specific behavior: You need to ensure that you can back up db1 to storage1.

Why A is Wrong: Generate a shared access signature (SAS) is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that you can back up db1 to storage1.

Why C is Wrong: Create an access policy. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: Rotate the storage keys is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that you can back up db1 to storage1.

質問: 140

14 TB の Microsoft SQL Server データベースを含むオンプレミス データセンターがあります。

Azure SQL マネージド インスタンスを作成し、オンプレミス データベースを新しいインスタンスに移行する予定です。

SQL マネージドインスタンスをサポートする 3 つのサービス レベルはどれですか。正解ごとに完全なソリューションが提示されます。注: 正解を選択するごとに 1 ポイントが加算されます。

- A. 汎用標準
- B. ビジネスクリティカルプレミアム
- C. ビジネスクリティカルメモリ最適化プレミアム
- D. 汎用プレミアム
- E. ビジネスクリティカルスタンダード

正解: ([正解を表示します](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why A and C and D is Correct: General Purpose Standard is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. In this scenario, the important constraint is: Which three service tiers support the SQL managed instance?

General Purpose Standard satisfies that constraint without adding an unrelated service or manual process.

Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. This fits the case because the requested outcome is: Which three service tiers support the SQL managed instance? The selected feature addresses that outcome directly rather than relying on a workaround. General Purpose Premium is correct because it is the feature whose normal purpose matches the stated requirement. General Purpose Premium is an Azure data- platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario wording points to that specific behavior: Which three service tiers support the SQL managed instance?

Why B is Wrong: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why E is Wrong: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

質問: 141

Azureサブスクリプションをお持ちです。このサブスクリプションには、Azure Virtual Machines上のSQL Serverインスタンス(SQL1)とAzure Automationアカウント(account1)が含まれています。SQL Serverエージェントサービスが停止した場合に再起動するように、account1を構成する必要があります。どのような設定を行う必要がありますか？

- A. VMの起動/停止
- B. 更新管理
- C. 状態構成 (DSC)
- D. 変更追跡

正解: C ([コメントを發表する](#))

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why C is Correct: State configuration (DSC) can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. This fits the case because the requested outcome is: You need to configure account1 to restart the SQL Server Agent service if the service stops. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Start/Stop VM can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

Why B is Wrong: Update management can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It would leave too much manual work or would operate at the wrong scope for the requested automation.

Why D is Wrong: Change tracking can participate in administration workflows, but automation questions require the mechanism that can execute repeatable tasks with the correct scope and schedule. It belongs to a different step of operations; this question needs the mechanism that can run or deploy the task repeatedly.

質問: 142

SERVER1データベースの移行を計画しています。ソリューションはビジネス要件を満たす必要があります。

移行計画には何を含めるべきですか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。




Azure Database Migration Service pricing tier: ▼

- Standard 2-vCore
- Standard 4-vCore
- Premium 4-vCore

Required Azure resource: ▼

- A virtual network that has service endpoints
- A VPN gateway
- An Azure Logic app

正解:



Azure Database Migration Service pricing tier: ▼


- Standard 2-vCore
- Standard 4-vCore
- Premium 4-vCore

Required Azure resource: ▼

- A virtual network that has service endpoints
- A VPN gateway
- An Azure Logic app

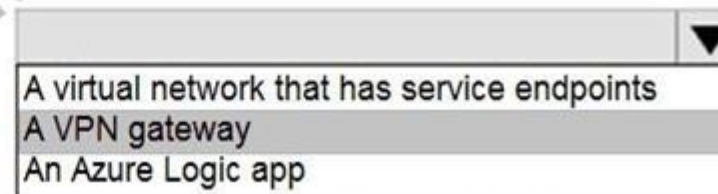
Explanation:

Azure Database Migration Service pricing tier:



Standard 2-vCore
Standard 4-vCore
Premium 4-vCore

Required Azure resource:



A virtual network that has service endpoints
A VPN gateway
An Azure Logic app

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.

Reference:

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

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: The solution must meet the business requirements.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

質問: 143

次のスクリプトを使用して Instance1 をデプロイする予定です。

```
{
  "type": "Microsoft.Sql/managedInstances",
  "apiVersion": "2019-06-01-preview",
  "name": "[parameters('instanceName')]",
  "location": "[resourceGroup().location]",
  "sku": {
    "name": "[parameters('skuName')]",
    "tier": "[parameters('skuEdition')]"
  },
  "dependsOn": [
    "Microsoft.Resources/deployments/BuildMInetworking"
  ],
  "properties": {
    "administratorLogin": "[parameters('adminLogin')]",
    "administratorLoginPassword": "[parameters('adminPassword')]",
    "subnetId": "[resourceId('Microsoft.Network/virtualNetworks/subnets', parameters('netName'), parameters('subnetName'))]",
    "licenseType": "OPEL",
    "storageRedundancy": "Geo"
  }
}
```

licenseType および storagedundancy パラメータを指定する必要があります。デプロイメントは、DB1 および DB2 の可用性要件とビジネス要件を満たす必要があります。各パラメータをどのように設定すればよいですか？ 回答するには、回答領域で適切なオプションを選択してください。

Answer Area

licenseType:

storageRedundancy:

Microsoft

正解:

Answer Area

licenseType:

storageRedundancy:

Microsoft

Explanation:

Answer Area

licenseType:

storageRedundancy:

Microsoft

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to specify the licenseType and storagedundancy parameters.

Why the alternate choices are Wrong: The alternate choices do not satisfy the stated availability design because quorum and listener resources must match Azure failover-cluster behavior and the required SLA.

質問: 144

Azure SQL データベースには Customers という名前のテーブルがあり、Customers には次の列が含まれています。

Name	Data type
CustomerID	int
FirstName	nvarchar(50)
LastName	nvarchar(50)
Address	nvarchar(1024)
PhoneNumber	varchar(20)
Email	varchar(255)
CreatedDate	datetime
LastUpdateDate	datetime

電話番号とメールアドレスの列のデータは暗号化されています。

顧客向けのデータ分類戦略を提案する必要があります。提案するソリューションは、以下の要件を満たす必要があります。

* データセキュリティおよびプライバシー要件をサポートする。

機密データへのアクセス監査をサポートする

管理業務の手間を最小限に抑える。

何をおすすめしますか？

- A. 各列に感度ラベルを適用します。
- B. データベースに対して透過的データ暗号化 (TDE) を有効にします。
- C. 電話番号とメールアドレスの列に対して動的なデータマスキングを実装します。
- D. 顧客にセンシティブラベルを適用する

正解: A (コメントを发表する)

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why A is Correct: Apply a sensitivity label to each column. is correct because it is the feature whose normal purpose matches the stated requirement. Apply a sensitivity label to each column. is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario wording points to that specific behavior: You need to recommend a data classification strategy for Customers.

Why B is Wrong: Transparent Data Encryption protects database files, backups, and logs at rest; it does not encrypt data in use or enforce user-level data visibility. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to recommend a data classification strategy for Customers.

Why C is Wrong: Dynamic Data Masking obscures sensitive values in query results for non-privileged users, but it does not restrict which rows are returned. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why D is Wrong: Apply a sensitivity label to Customers is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

質問: 145

Azure Virtual Machines インスタンスに SQLVM1 という名前の SQL Server があり、このインスタンスには DB1 と DB2 という 2 つのデータベースが含まれています。DB1 と DB2 のデータベースとログファイルは、マネージド ディスク上にホストされています。

DB1とDB2のスナップショットバックアップを実行する必要があります

1 SQL ステートメントをどのように完了する必要がありますか? 回答するには、回答領域で適切なオプションを選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area

```

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = WORKLOAD GROUP);
BACKUP GROUP db1, db2
TO DISK = 'f:\backup\db.bkm'
WITH FILE_SNAPSHOT, FORMAT;
  
```

Microsoft

正解:

Answer Area

Microsoft

```

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = WORKLOAD GROUP);
BACKUP GROUP db1, db2
TO DISK = 'f:\backup\db.bkm'
WITH FILE_SNAPSHOT, FORMAT;
  
```

Explanation:

Answer Area

```

ALTER SERVER CONFIGURATION
SET SUSPEND_FOR_SNAPSHOT_BACKUP = ON
(GROUP = (db1, db2));
BACKUP GROUP db1, db2
TO DISK = 'f:\backup\db.bkm'
WITH FILE_SNAPSHOT, FORMAT;
  
```

Microsoft


Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to perform a snapshot backup of DB1 and DB2 How should you complete the T-SQL statements?

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 146

Table1 という名前のテーブルを含む DB1 という名前の Microsoft SQL Server データベースがあります。
User1 という名前のユーザーのデータベース ロール メンバーシップを次の図に示します。



ドロップダウン メニューを使用して、グラフィックに表示された情報に基づいて各ステートメントを完成させる回答の選択肢を選択します。
注意: 正しい選択ごとに 1 ポイントが加算されます。

User1 can [answer choice].

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

add a column to Table1
delete a row from Table1
delete Table1

db_datareader
db_ddladmin
db_denydatareader
db_denydatawriter

正解:

User1 can [answer choice].

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

add a column to Table1
delete a row from Table1
delete Table1

db_datareader
db_ddladmin
db_denydatareader
db_denydatawriter

Explanation:

User1 can [answer choice].

To ensure that User1 can run queries to retrieve data from DB1, you must assign User1 the [answer choice] database role.

add a column to Table1
delete a row from Table1
delete Table1

db_datareader
db_ddladmin
db_denydatareader
db_denydatawriter

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> Basic Concept: This question tests implement a secure environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have a Microsoft SQL Server database named DB1 that contains a table named Table1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 147

企業では、防犯カメラの画像を分析し、異常なアクティビティに対応するセキュリティチームにアラートを送信しています。このソリューションでは、Azure Databricks を使用しています。Apache Spark レベルのイベント、Spark Structured Streaming メトリック、およびアプリケーション メトリックを Azure Monitor に送信する必要があります。順番に実行する必要がある 3 つのアクションはどれですか。回答するには、回答エリアのアクション リストから適切なアクションを移動し、正しい順序に並べます。

Actions

Deploy Grafana to an Azure virtual machine.

Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.

Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.

Answer Area

正解:

Actions

Deploy Grafana to an Azure virtual machine.

Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.

Create Dropwizard counters in the application code.

Create a data source in Azure Monitor.

Configure the Databricks cluster to use the Databricks monitoring library.


Answer Area

Configure the Databricks cluster to use the Databricks monitoring library.

Build a **spark-listeners-loganalytics-1.0-SNAPSHOT.jar** JAR file.

Create Dropwizard counters in the application code.

Explanation:



Send application metrics using Dropwizard.

Spark uses a configurable metrics system based on the Dropwizard Metrics Library.

To send application metrics from Azure Databricks application code to Azure Monitor, follow these steps:

Step 1: Configure your Azure Databricks cluster to use the Databricksmonitoring library.

Prerequisite: Configure your Azure Databricks cluster to use the monitoring library.

Step 2: Build the spark-listeners-loganalytics-1.0-SNAPSHOT.jar JAR file Step 3: Create Dropwizard counters in your application code Create Dropwizard gauges or counters in your application code

質問: 148

Azure サブスクリプションがあります。

Azure SQL データベースをデプロイする必要があります。ソリューションは、次の要件を満たす必要があります。

* CPU リソースを動的にスケーリングします。

* コストを削減するために、データベースを一時停止できることを確認してください。

何を使うべきですか？

- A. ビジネス クリティカル サービス層
- B. サーバーレス コンピューティング レベル
- C. エラスティック プール
- D. 汎用サービス層

正解: (正解を表示します)

Basic Concept: This question tests Azure monitoring and performance management, including which metric, log destination, or alerting behavior fits the operational requirement.

Why B is Correct: the serverless compute tier is correct because it is the feature whose normal purpose matches the stated requirement. The serverless tier automatically scales compute and can pause during inactive periods, making it suitable for intermittent workloads with bursts and idle windows. The scenario wording points to that specific behavior: You need to deploy an Azure SQL database.

Why A is Wrong: Business Critical provides low-latency local SSD storage and high availability for demanding workloads, but it is usually more expensive than necessary for bursty idle workloads. That design keeps capacity available for a different kind of workload and does not match the cost behavior expected from an idle-and-burst database pattern.

Why C is Wrong: Elastic pools let multiple Azure SQL databases share a fixed pool of resources, which is cost-effective when individual database usage is variable and unpredictable. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why D is Wrong: the General Purpose service tier is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

質問: 149

Server 1 という名前の Azure SQL マネージド インスタンスを含む Azure サブスクリプションがあります。

サーバーの監査を有効にする必要があります。ソリューションは管理作業を最小限に抑える必要があります。

どの 4 つのアクションを順番に実行する必要がありますか? 回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions

- On Server1, enable diagnostics.
- On Server1, run the create server audit specification T-SQL statement.
- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, run the create server audit T-SQL statement.

Answer Area

正解:

Microsoft

Actions

- On Server1, enable diagnostics.
- On Server1, run the create server audit specification T-SQL statement.
- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, run the create server audit T-SQL statement.

Answer Area

- From the Azure portal, create an Azure Storage account.
- From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account.
- On Server1, create a credential.
- On Server1, run the create server audit T-SQL statement.

Explanation:

Actions	Answer Area
⋮ On Server1, enable diagnostics. ⋮ On Server1, run the create server audit specification T-SQL statement.	1 ⋮ From the Azure portal, create an Azure Storage account. 2 ⋮ From the Azure portal, copy the URL and the shared access signature (SAS) token of the storage account. 3 ⋮ On Server1, create a credential. 4 ⋮ On Server1, run the create server audit T-SQL statement.

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to enable auditing for Server!.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 150

VM1 と VM2 という名前の Azure Virtual Machines 上の SQL Server の 2 つのインスタンスを含む Azure サブスクリプションがあります。どちらのインスタンスも Microsoft SQL Server 2019 CU8 を実行します。フェールオーバー クラスター インスタンス (FCI) を VM1 と VM2 にデプロイする必要があります。ソリューションでは、次の必要性を排除する必要があります。

* 分散ネットワーク名 (DNN)

* ロードバランサー

あなたは何をするべきか？

- A. VM1 と VM2 を単一の近接配置グループに展開します。
- B. VM1 と VM2 を同じ Azure リージョン内の異なる近接配置グループにデプロイします。
- C. VM1 と VM2 を単一のサブネットに接続します。
- D. VM1 と VM2 を単一の仮想ネットワーク上の異なるサブネットに接続します。

正解: ([正解を表示します](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why D is Correct: Connect VM1 and VM2 to different subnets on a single virtual network. is correct because it is the feature whose normal purpose matches the stated requirement. Connect VM1 and VM2 to different subnets on a single virtual network. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario wording points to that specific behavior: You need to deploy a failover cluster instance (FCI) to VM1 and VM2.

Why A is Wrong: Deploy VM1 and VM2 to a single proximity placement group. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why B is Wrong: Deploy VM1 and VM2 to different proximity placement groups in the same Azure region.

is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: Connect VM1 and VM2 to a single subnet. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: 151

Azure仮想マシンでホストされているSQLServerデータベースのデータベースバックアップソリューションを構築しています。

Azureの地域が停止した場合は、データベースのバックアップを復元する必要があります。ソリューションはコストを最小限に抑える必要があります。

バックアップにはどのタイプのストレージアカウントを使用する必要がありますか？

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

正解: B ([コメントを发表する](#))

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).

Reference:

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

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> 503問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 152

DB1 という名前の Azure SQL データベースがあります。DB1 には、Col1 という名前の列を持つテーブルが含まれています。

Col1 のデータを暗号化する必要があります。

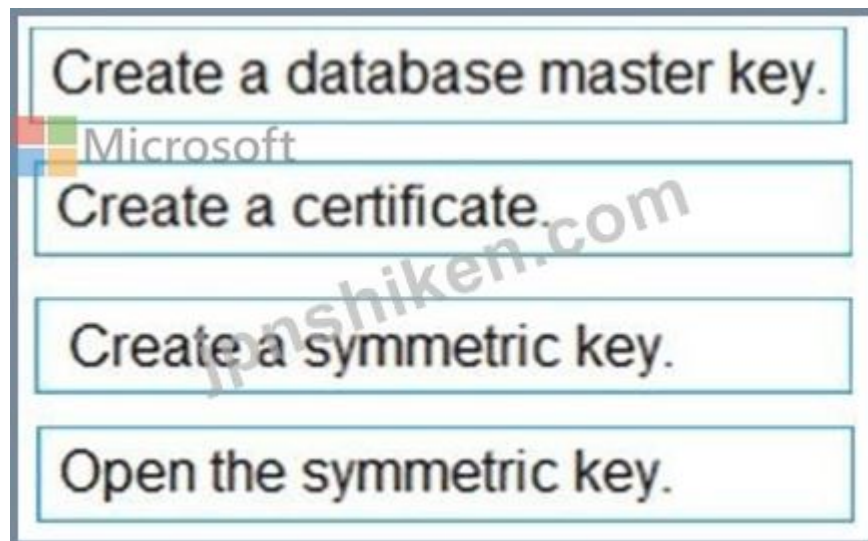
DB1 に対して順番に実行する必要がある 4 つのアクションはどれですか。回答するには、適切なアクションをアクション リストから回答領域に移動し、正しい順序に並べます。

Actions	Answer Area
Create a database master key.	
Create a column master key.	
Open the symmetric key.	
Create a certificate.	
Update Col1.	
Create a symmetric key.	

正解:

Actions	Answer Area
Create a database master key.	Create a database master key.
Create a column master key.	Create a certificate.
Open the symmetric key.	Create a symmetric key.
Create a certificate.	Open the symmetric key.
Update Col1.	
Create a symmetric key.	

Explanation:



Create a database master key.

Microsoft

Create a certificate.

Create a symmetric key.

Open the symmetric key.

Reference:

<https://www.sqlshack.com/an-overview-of-the-column-level-sql-server-encryption/> Basic Concept: This question tests implement a secure environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to encrypt the data in Col1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 153

Server1 という名前の Azure SQL Database の論理サーバーがあり、このサーバーには Pool1 という名前のエラスティック プールがあります。Watcher 1 という名前のデータベース ウォッチャーがあります。次の KQL クエリを実行します。



```
let logicalServer = @"Server1";
let elasticPoolName = @"Pool1";
let replicaType = "Primary";
let duration = 20h;
sqldb_elastic_pool_resource_utilization
| where sample_time_utc > ago(duration)
| where logical_server_name == logicalServer
| where replica_type == replicaType
| project sample_time_utc,
    avg_cpu_percent,
    avg_instance_cpu_percent,
    avg_data_io_percent,
    avg_log_write_percent,
    max_worker_percent
| sort by sample_time_utc desc
| render timechart;
```

以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Answer Area



Microsoft

Statements

The query returns the execution time of individual queries.

Yes

No

☐☐

The query returns the resource consumption metrics from the last day.

☐☐

The query visualizes the result set and includes a legend for each series.

☐☐

正解:

Answer Area



Microsoft

Statements

The query returns the execution time of individual queries.

Yes

No

☐☐

The query returns the resource consumption metrics from the last day.

☒☐

The query visualizes the result set and includes a legend for each series.

☒☐

Explanation:

Answer Area



Statements

The query returns the execution time of individual queries.

Yes

No

☐☒

The query returns the resource consumption metrics from the last day.

☒☐

The query visualizes the result set and includes a legend for each series.

☒☐

Basic Concept: This question tests Azure SQL sizing and purchasing choices, where compute model, scaling behavior, licensing benefits, and cost pattern must be evaluated together.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have a logical server for Azure SQL Database named Server1 that has an elastic pool named Pool1. You have a database watcher named Watcher 1.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 154

オンプレミスのMicrosoft SQL Server 2022インスタンス(SQL1)があります。Azureサブスクリプションには、Azure SQL Managed Instance(SQLMI1)が含まれています。SQL1とSQLMHの間で双方向のディザスターリカバリーソリューションを構成する必要があります。ソリューションには何を含めるべきですか？

- A. ログ再生サービス (LRS)
- B. マネージドインスタンスリンク
- C. Azure データベース移行サービス
- D. フェイルオーバーグループ

正解: **B** ([コメントを发表する](#))

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why B is Correct: the Managed instance link is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement.

This fits the case because the requested outcome is: You need to configure a bidirectional disaster recovery solution between SQL1 and SQLMH. The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: Log Replay Service replays full, differential, and transaction log backups into Azure SQL Managed Instance, making it appropriate for low-downtime migrations from SQL Server. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why C is Wrong: Azure Database Migration Service is the managed migration service used to move database engines to Azure targets with online or offline migration modes depending on source and target support. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why D is Wrong: Auto-failover groups provide managed geo-replication and listener-based failover for Azure SQL Database or Managed Instance, including automatic failover policies for regional outages. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: **155**

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、述べられた目標を達成する可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

ステージングゾーンを含むAzureData LakeStorageアカウントがあります。

ステージングゾーンから増分データを取り込み、Rスクリプトを実行してデータを変換し、変換されたデータをAzure SynapseAnalyticsのデータウェアハウスに挿入する毎日のプロセスを設計する必要があります。

解決策 :Azure Data Factoryスケジュールトリガーを使用して、マッピングデータフローを実行するパイプラインを実行してから、データウェアハウスにデータを挿入します。

これは目標を達成していますか？

- A. はい
- B. いいえ

正解: ([正解を表示します](#))

If you need to transform data in a way that is not supported by Data Factory, you can create a custom activity, not a mapping flow,5 with your own data processing logic and use the activity in the pipeline. You can create a custom activity to run R scripts on your HDInsight cluster with R installed.

Reference:

<https://docs.microsoft.com/en-US/azure/data-factory/transform-data>

Basic Concept: This question tests task automation for database administration, including which Azure or SQL mechanism can run repeatable maintenance and configuration work.

Why B is Correct: No is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario asks for: Each question in the series contains a unique solution that might meet the stated goals. That makes No the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It does not expose the required metric, query history, wait/blocking signal, or tuning mechanism; using it would not give the administrator the evidence requested.

質問: **156**

DB1という名前のAzure SQLデータベースがあり、厳格な監査要件が設定されています。

DB1 で台帳機能を有効にし、自動ダイジェストストレージを使用してデータベース検証を実装する予定です。検証用ストアードプロシージャを実行する前に、DB1 のデータベース分離レベルを設定する必要があります。

隔離レベルはどのくらいに設定すべきですか？

A. 読み取りがコミットされました

B. 未コミットの読み取り

C. スナップショット

D. 繰り返し読み取り可能

正解: ([正解を表示します](#))

Basic Concept: This question tests database auditing and query visibility. DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why C is Correct: Snapshot is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The scenario asks for: You have an Azure SQL database named DB1 that has strict auditing requirements. That makes Snapshot the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Read committed is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You have an Azure SQL database named DB1 that has strict auditing requirements.

Why B is Wrong: Read uncommitted is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It protects or manages a different security layer, so the required database access or protection behavior would still be incomplete.

Why D is Wrong: Repeatable read is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You have an Azure SQL database named DB1 that has strict auditing requirements.

質問: **157**

標準料金階層にworkspace1という名前のAzureDatabricksワークスペースがあります。Workspace1には、cluster1という名前の汎用クラスターが含まれています。

cluster1の起動とスケールアップにかかる時間を短縮する必要があります。ソリューションはコストを最小限に抑える必要があります。

あなたは最初に何をすべきですか？

A. workspace1をプレミアム料金階層にアップグレードします。

B. workspace1のグローバル初期化スクリプトを構成します。

C. workspace1にプールを作成します。

D. workspace1でクラスターポリシーを作成します。

正解: ([正解を表示します](#))

You can use Databricks Pools to Speed up your Data Pipelines and Scale Clusters Quickly.

Databricks Pools, a managed cache of virtual machine instances that enables clusters to start and scale 4 times faster.

Reference:

<https://databricks.com/blog/2019/11/11/databricks-pools-speed-up-data-pipelines.html> Basic Concept: This question tests plan and configure a high availability and disaster recovery (ha/dr) environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: Create a pool in workspace1. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. The scenario asks for: You need to reduce the time it takes for cluster1 to start and scale up. That makes Create a pool in workspace1. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Upgrade workspace1 to the Premium pricing tier. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

Why B is Wrong: Configure a global init script for workspace1. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: Create a cluster policy in workspace1. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: 158

DB1という名前のデータベースを含むAzure仮想マシン上にSQLServerがあります。

次の図に示すように、クエリ1178902の各実行の期間をミリ秒単位で示すプランの概要を表示します。



クエリが最短時間で実行される実行プランを使用するには、どうすればよいですか？

- A. プラン1221065のクエリ実行プランを強制します。
- B. DBCCFREEPROCCACHEコマンドを実行します。
- C. プラン1220917のクエリ実行プランを強制します。
- D. パラメータスニффイングを無効にします。

正解: ([正解を表示します](#))

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/query-store-usage-scenarios> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why C is Correct: Force the query execution plan for plan 1220917. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario asks for: 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? That makes Force the query execution plan for plan 1220917. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Force the query execution plan for plan 1221065. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

Why B is Wrong: Run the DBCC FREEPROCCACHE command. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

Why D is Wrong: Disable parameter sniffing. is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It is useful in other troubleshooting paths, but this scenario requires a more specific monitoring or optimization feature.

質問: **159**

タスク5

db1のCPU使用率が1分間隔でサンプリングされた5分間の平均で50%を超えた場合、メールアラートを生成する必要があります。アラートはadmin@contoso.comに送信してください。

SQL Server Management StudioとAzureポータルを使用する必要があるかもしれません。

正解:

See the explanation part for the complete Solution.

Explanation:

Create an Azure Monitor metric alert rule on database db1 using the metric:

CPU percentage

Configure the condition as:

Aggregation: Average

Operator: Greater than

Threshold: 50

Aggregation granularity / Period: 1 minute

Frequency of evaluation: 1 minute

Evaluation period / Lookback window: 5 minutes

Action group email: admin@contoso.com

Azure SQL Database exposes CPU percentage as a platform metric, and Azure Monitor metric alerts can send notifications through action groups such as email. Microsoft describes SQL Database alerts as metric-based alerts that can send email when metrics such as CPU usage reach a defined threshold.

Azure Portal Method - Recommended for Simulation

Step 1: Open the db1 Azure SQL database

Sign in to the Azure portal.

Search for SQL databases.

Open the database named:

db1

Do not open the SQL logical server unless the alert needs to apply to all databases. This task is specifically for db1, so the alert scope must be the db1 database resource.

Step 2: Create a new alert rule

From the db1 database page:

In the left menu, select Alerts.

Select Create.

Select Alert rule.

Microsoft's Azure Monitor workflow allows you to create an alert rule directly from the target resource.

When you create it from a resource, the resource is automatically set as the alert scope.

Step 3: Confirm the alert scope

On the alert rule page, confirm the scope is the Azure SQL database:

db1

The resource type should be similar to:

SQL database

Microsoft.Sql/servers/databases

If the scope is the SQL server instead of the database, remove it and select the db1 database resource.

Step 4: Add the alert condition

Under Condition, select Add condition.

Choose the metric:

CPU percentage

Sometimes it appears as:

CPU percent

or metric name:

cpu_percent

For Azure SQL Database, CPU percentage represents CPU consumption toward the database workload limit, expressed as a percentage.

Step 5: Configure the signal logic

Configure the alert logic exactly like this:

Setting

Value

Threshold type

Static

Aggregation type

Average

Operator

Greater than

Threshold value

50

Unit

Percent

Aggregation granularity / Period

1 minute

Frequency of evaluation

1 minute

Evaluation period / Lookback window

5 minutes

This means Azure Monitor evaluates the CPU metric every minute, using one-minute metric samples, and fires the alert only when the average CPU percentage is greater than 50% across the five-minute evaluation window.

Be precise here. The task says:

average CPU percentage utilization is greater than 50 percent for five minutes sampled at one-minute intervals So the correct choices are:

Average

Greater than 50

Every 1 minute

Over the last 5 minutes

Azure Monitor metric alert rules combine the monitored resource, the metric condition, and action groups that run when the condition is met.

Step 6: Create the email action group

Under Actions, select:

Create action group

Configure the basics:

Setting

Value

Subscription

Use the current subscription

Resource group

Use the lab resource group

Action group name

AG-db1-CPU-Email

Display name

db1CPU

Then go to Notifications.

Add a notification:

Setting

Value

Notification type

Email/SMS message/Push/Voice

Name

EmailAdmin

Email

admin@contoso.com

Select OK, then Review + create, then Create.

Action groups define the notification or automation action that runs when an alert fires. Microsoft documents email as a supported action group notification type.

Step 7: Configure alert rule details

Under Details, configure:

Setting

Value

Severity

2 or 3

Alert rule name

db1 CPU greater than 50 percent

Description

Alert when average CPU percentage for db1 is greater than 50 percent for 5 minutes.

Enable upon creation

Yes

Severity is usually not specified by the task, so any reasonable severity is acceptable. In an exam lab, I would use Severity 2 for CPU performance impact.

Step 8: Review and create

Select Review + create.

Confirm the condition says something equivalent to:

Whenever the average CPU percentage is greater than 50

Confirm the evaluation settings are:

Check every 1 minute

Lookback period 5 minutes

Confirm the action group sends email to:

admin@contoso.com

Select Create.

Verification

After creation:

Open db1.

Go to Alerts.

Select Alert rules.

Confirm the alert rule exists and is enabled.

Open the rule and verify:

Scope: db1

Signal: CPU percentage

Aggregation: Average

Operator: Greater than

Threshold: 50

Evaluation frequency: 1 minute

Window size: 5 minutes

Action group: admin@contoso.com

Azure CLI Method

Use this only if the simulation gives you Cloud Shell.

First get the database resource ID:

```
az sql db show \
```

```
--resource-group < resource-group-name > \
```

```
--server < sql-server-name > \
```

```
--name db1 \
```

```
--query id \
```

```
--output tsv
```

Create the action group:

```
az monitor action-group create \
```

```
--resource-group < resource-group-name > \
```

```
--name AG-db1-CPU-Email \
```

```
--short-name db1CPU \
```

```
--action email EmailAdmin admin@contoso.com
```

Create the metric alert:

```
az monitor metrics alert create \
```

```
--name " db1 CPU greater than 50 percent " \
```

```
--resource-group < resource-group-name > \
```

```
--scopes < db1-resource-id > \
```

```
--condition " avg cpu_percent > 50 " \
```

--window-size 5m \
--evaluation-frequency 1m \
--action AG-db1-CPU-Email \
--description " Alert when average CPU percentage for db1 is greater than 50 percent for 5 minutes. " The metric name commonly used for Azure SQL Database CPU percentage in CLI/ARM contexts is:

cpu_percent

SSMS Clarification

SSMS is not the correct tool for this task.

Do not configure Database Mail. Azure SQL Database does not use SQL Server Agent/Database Mail in the same way as SQL Server on a VM or SQL Managed Instance. This requirement is an Azure Monitor metric alert requirement.

Correct tool:

Azure portal > db1 > Alerts > Create alert rule

Wrong tool:

SSMS Database Mail

Final Exam-Lab Configuration

Use this exact configuration:

Resource: db1

Alert type: Metric alert

Metric: CPU percentage

Aggregation: Average

Operator: Greater than

Threshold: 50

Aggregation granularity: 1 minute

Evaluation frequency: 1 minute

Evaluation period/window: 5 minutes

Action group notification: Email

Email recipient: admin@contoso.com

That completes the task.

質問: 160

オンプレミスのMicrosoftSQLServerデータベースをAzureSQLDatabaseに移行する必要があります。ソリューションはダウンタイムを最小限に抑える必要があります。

あなたは何をするべきか？

- A. トランザクションログ配布を構成します。
- B. AlwaysOn可用性グループを実装します。
- C. トランザクションレプリケーションを構成します。
- D. BACPACをインポートします。

正解: C ([コメントを發表する](#))

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/migrate-to-database-from-sql-server#method-1- migration-with-downtime-during-the-migration> Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why C is Correct: Transactional replication distributes committed changes to subscribers, but it is not the default full-database migration method for Azure SQL Managed Instance when the requirement is minimal effort. The scenario asks for: You need to migrate an on-premises Microsoft SQL Server database to Azure SQL Database. That makes Configure transactional replication. the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Log shipping is a SQL Server HA/DR technique that copies and restores log backups to a secondary database; it is more manual and operationally heavier for Azure SQL Managed Instance migration than LRS. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.


Why B is Wrong: Always On availability groups are SQL Server/SQL Managed Instance HA technology, not the Azure SQL Database PaaS failover abstraction used for database groups. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: Import a BACPAC. is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It does not meet the failover, restore, quorum, or cross-region continuity target stated in the question, even if it is valid in a different availability design.

質問: **161**

Azure 仮想マシン上に SQL Server があります。

次の図に示すクエリ プランを確認します。



以下の各文について、正しい場合は「はい」を選択してください。そうでない場合は「いいえ」を選択してください。

注意: 正しい選択ごとに 1 ポイントが加算されます。

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	☐	☐
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	☐	☐
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.	☐	☐

正解:

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	☒	☐
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	☐	☒
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.	☒	☐

Explanation:

Statements	Yes	No
You will reduce the I/O usage and the query execution time if you force the query plan.	☒	☐
You will increase the I/O usage and the query execution time if you create a new index on the SalesOrderHeader table.	☐	☒
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.	☒	☐

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-the-query-store> Basic Concept: This question tests database auditing and query visibility.

DP-300 expects you to distinguish between collecting audit events, retaining logs, and analyzing query behavior.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You have SQL Server on an Azure virtual machine.

Why the alternate choices are Wrong: The alternate selections apply to different administrative stages or different scopes. In a hotspot or drag-drop task, order and scope are part of the answer, so a technically valid Azure feature can still be wrong if it is placed in the wrong step.

質問: 162

DB1という名前のオンプレミスMicrosoftSQL Server2016データベースにデータを格納するApp1という名前のオンプレミスアプリがあります。

App1の追加のインスタンスをデプロイしてAzureリージョンを分離することを計画しています。各リージョンには、App1とDB1の個別のインスタンスがあります。DB1の個別のインスタンスは、Azure SQL DataSyncを使用して同期します。

展開にはデータベースサービスを推奨する必要があります。このソリューションでは、管理作業を最小限に抑える必要があります。

推奨事項には何を含める必要がありますか？

- A. AzureSQLマネージドインスタンス
- B. AzureSQLデータベースの単一データベース
- C. PostgreSQL用のAzureデータベース
- D. Azure仮想マシン上のSQL Server

正解: ([正解を表示します](#))

Azure SQL Database single database supports Data Sync.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/features-comparison> Basic Concept: This question tests plan and implement data platform resources in the DP-300 exam context.

The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: Azure SQL Database single database is correct because it is the feature whose normal purpose matches the stated requirement. Azure SQL Database single database is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. The scenario wording points to that specific behavior: You need to recommend a database service for the deployment.

Why A is Wrong: Azure SQL Managed instance is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why C is Wrong: Azure Database for PostgreSQL is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It applies to a different Azure data-platform design than the one the question is testing.

Why D is Wrong: SQL Server on Azure virtual machines is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

質問: **163**

DB1という名前のデータベースを含むAzure仮想マシン上にSQLServerがあります。DB1には、CustomerPIIという名前のテーブルが含まれています。

ユーザーがCustomerPIIテーブルをクエリするたびに記録する必要があります。

どの2つのオプションを有効にする必要がありますか？それぞれの正解は、解決策の一部を示しています。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. サーバー監査仕様
- B. SQLServer監査
- C. データベース監査仕様
- D. サーバープリンシパル

正解: ([正解を表示します](#))

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>

質問: **164**

PostgreSQL データベースを移行するには何を使用すればよいですか？

- A. Azure データ ボックス
- B. AzCopy
- C. Azure データベース移行サービス
- D. Azure サイトの回復

正解: **C** ([コメントを发表する](#))

Reference:

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

Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why C is Correct: Azure Database Migration Service is the managed migration service used to move database engines to Azure targets with online or offline migration modes depending on source and target support. The scenario asks for: What should you use to migrate the PostgreSQL database? That makes Azure Database Migration Service the option that satisfies the required Azure SQL layer and operational outcome.

Why A is Wrong: Azure Data Box is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

Why B is Wrong: AzCopy is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It mismatches the target platform, migration method, compute tier, or deployment model described in the scenario, so it would not satisfy the workload as stated.

Why D is Wrong: Azure Site Recovery is an Azure data-platform feature, but it must match the workload model, service tier, migration path, or deployment constraint described in the scenario. It would solve a neighboring problem, but not the deployment, sizing, or migration requirement in this item.

質問: **165**

Azure サブスクリプションをお持ちです

Azure仮想マシン上にSQL Serverのインスタンスを作成する予定です。データセンターの障害に対してインスタンスが耐障害性を持つようにする必要があります。何を使用すればよいでしょうか？

A. 可用性セット

B. Azure Site Recovery

C. 可用性ゾーン

D. 近接配置グループ

正解: ([正解を表示します](#))

Basic Concept: This question tests high availability and disaster recovery design for Azure SQL, SQL Server on Azure VMs, and regional failure scenarios.

Why B is Correct: Azure Site Recovery is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. This fits the case because the requested outcome is: You need to ensure that the instance is resilient against datacenter failures.

The selected feature addresses that outcome directly rather than relying on a workaround.

Why A is Wrong: an availability set is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

Why C is Wrong: an availability zone is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It handles a different resilience pattern and would not deliver the failover or recovery behavior required here.

Why D is Wrong: a proximity placement group is part of the availability or recovery design space, but the correct choice must satisfy the specified failover, restore, quorum, RPO, or RTO requirement. It is not wrong technology in general, but it is the wrong HA/DR control for this scenario ' s failure model.

質問: **166**

注 :この質問は、同じシナリオを提示する一連の質問の一部です。シリーズの各質問には、述べられた目標を達成する可能性のある独自の解決策が含まれています。一部の質問セットには複数の正しい解決策がある場合がありますが、他の質問セットには正しい解決策がない場合があります。

このセクションの質問に回答した後は、その質問に戻ることはできません。その結果、これらの質問はレビュー画面に表示されません。

Table1という名前のテーブルを含むAzureSynapseAnalytics専用のSQLプールがあります。

container1という名前のAzureData Lake StorageGen2コンテナに取り込まれてロードされるファイルがあります。

ファイルからTable1にデータを挿入し、データを変換することを計画しています。ファイル内のデータの各行は、Table1のサービングレイヤーに1つの行を生成します。

ソースデータファイルがcontainer1にロードされるときに、DateTimeがTable1の追加の列として格納されていることを確認する必要があります。

解決策：専用SQLプールを使用して、追加のDateTime列を持つ外部テーブルを作成します。

これは目標を達成していますか？

A. はい

B. いいえ

正解: B ([コメントを发表する](#))

Instead use a serverless SQL pool to create an external table with the extra column.

Note: In dedicated SQL pools you can only use Parquet native external tables. Native external tables are generally available in serverless SQL pools.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/create-use-external-tables> Basic Concept: This question tests monitor, configure, and optimize database resources in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: No is correct because it is the feature whose normal purpose matches the stated requirement. No is related to operational monitoring or tuning, but it must match the exact signal needed:

query history, resource utilization, wait/blocking detail, or automatic remediation. The scenario wording points to that specific behavior: Each question in the series contains a unique solution that might meet the stated goals.

Why A is Wrong: Yes is related to operational monitoring or tuning, but it must match the exact signal needed: query history, resource utilization, wait/blocking detail, or automatic remediation. It would produce a different operational signal than the one needed to investigate, alert, or tune the workload in this question.

有効的な**DP-300J**問題集はJPNTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTest.comは今最新**DP-300J**試験問題集を提供します。JPNTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503問**、**30%ディスカウント**、特別な割引コード: **JPNshiken**」

質問: 167

タスク11

Microsoft SQL Server 2012 用に作成されたレガシー アプリケーションがあります。このアプリケーションは、db1 にアクセスする唯一のアプリケーションになります。db1 が SQL Server 2012 のすべての機能および構文と互換性があることを確認する必要があります。

正解:

See the explanation part for the complete Solution.

Explanation:

To ensure that db1 is compatible with all the features and syntax of SQL Server 2012, you need to set the compatibility level of the database to 110, which is the compatibility level for SQL Server 2012. The compatibility level affects the behavior of certain Transact-SQL statements and features, and determines how the database engine interprets the SQL code.

You can set the compatibility level of db1 by using the Azure portal or Transact-SQL statements. Here are the steps for both methods:

Using the Azure portal:

Go to the Azure portal and select your Azure SQL Database server that hosts db1.

Select the database db1 and click on Query Performance Insight in the left menu.

Click on Configure Query Store and select 110 from the Compatibility level dropdown list.

Click on Save to apply the change.

Using Transact-SQL statements:

Connect to db1 using SQL Server Management Studio, Azure Data Studio, or any other tool that supports Transact-SQL statements.

Open a new query window and run the following command: ALTER DATABASE db1 SET COMPATIBILITY_LEVEL = 110; GO This command will set the compatibility level of db1 to 110, which is equivalent to SQL Server 2012.

These are the steps to set the compatibility level of db1 to 110.

質問: **168**

タスク 10

sql37006S95 上のすべてのデータベースを SQL インジェクション攻撃から保護する必要があります。

正解:

See the explanation part for the complete Solution

Explanation:

SQL injection attacks are a type of cyberattack that exploit a vulnerability in the application code that interacts with the database. An attacker can inject malicious SQL statements into the user input, such as a form field or a URL parameter, and execute them on the database server, resulting in data theft, corruption, or unauthorized access¹.

To protect all the databases on sql37006S95 from SQL injection attacks, you need to follow some best practices for securing your application and database layers. Here are some of the recommended steps:

Use parameterized queries or stored procedures to separate the SQL code from the user input. This will prevent the user input from being interpreted as part of the SQL statement and avoid SQL injection²³.

Validate and sanitize the user input before passing it to the database. This will ensure that the input conforms to the expected format and type, and remove any potentially harmful characters or keywords⁴.

Implement least privilege access for the database users and roles. This will limit the permissions and actions that the application can perform on the database, and reduce the impact of a successful SQL injection attack⁵.

Enable Advanced Threat Protection for Azure SQL Database. This is a feature that detects and alerts you of anomalous activities and potential threats on your database, such as SQL injection, brute force attacks, or unusual access patterns. You can configure the alert settings and notifications using the Azure portal or PowerShell.

These are some of the steps to protect all the databases on sql37006S95 from SQL injection attacks.

質問: **169**

AzureSQLデータベースのAzureMonitorから多数のアラートを受け取ります。

アラートの数を減らす必要があります。アラートを受信する必要があるのは、長期間にわたって使用パターンに大幅な変更があった場合のみです。

どの2つのアクションを実行する必要がありますか？それぞれの正解は、解決策の一部を示しています。

注：正しい選択はそれぞれ1ポイントの価値があります。

- A. しきい値感度を高に設定
- B. アラートロジックのしきい値を動的に設定します
- C. アラートロジックのしきい値を静的に設定します
- D. しきい値感度を低に設定
- E. フォースプランをオンに設定

正解: ([正解を表示します](#))

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.

Reference:

質問: 170

Azureサブスクリプションに、Server1という論理SQLサーバーが含まれています。Server1のマスターデータベースには、User1というユーザーが含まれています。User1がServer1上にデータベースを作成できるようにする必要があります。User1にはどのデータベースロールを割り当てる必要がありますか？

- A. db_owner
- B. dbmanager
- C. dbo
- D. db_ddladmin

正解: ([正解を表示します](#))

Basic Concept: This question tests implement a secure environment in the DP-300 exam context. The correct choice is determined by the exact service boundary and operational requirement stated in the scenario.

Why B is Correct: dbmanager matches the expected DP-300 administration action. dbmanager is a security- related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. The question is not asking for a general Azure capability; it is asking for the feature that produces this result: You need to ensure that User1 can create databases on Server1.

Why A is Wrong: db_owner is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that User1 can create databases on Server1.

Why C is Wrong: dbo is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It is not the right enforcement point for this case; the scenario needs the control that governs the data or identity path being tested.

Why D is Wrong: db_ddladmin is a security-related control, but its value depends on whether the requirement is identity, encryption, auditing, network isolation, or data exposure reduction. It does not satisfy the required identity, encryption, firewall, auditing, or data-exposure boundary described in the scenario: You need to ensure that User1 can create databases on Server1.

質問: 171

S0LSVR1という名前のオンプレミスのMicrosoft SQL Serverインスタンスがあり、その中にDB1という名前のデータベースがホストされています。

お客様は、SQLDB1 という名前の Azure SQL データベースを含む Azure サブスクリプションをお持ちです。


SQL Serverレプリケーションを使用して、DB1に保存されているデータをSQLDB1に移行する必要があります。このソリューションは、DB1のパフォーマンスへの影響を最小限に抑える必要があります。

レプリケーションはどのように構成すればよいですか？回答するには、回答欄で適切なオプションを選択してください。

注：正解ごとに1ポイントが加算されます。



正解:



Explanation:



Basic Concept: This question tests choosing the correct Azure migration approach by matching the source platform, target service, downtime tolerance, and administrative effort.

Why the selected answer is Correct: The selected values fit this scenario because they apply the required Azure SQL configuration at the correct scope and sequence: You need to migrate the data stored in DB1 to SQLDB1 by using SQL Server replication.

Why the alternate choices are Wrong: The alternate selections would either use a migration tier that does not match the downtime requirement, choose a network resource that does not support the migration path, or add an unnecessary deployment component.

有効的な**DP-300J**問題集はJPNTTest.com提供され、**DP-300J**試験に合格することに役に立ちます！JPNTTest.comは今最新**DP-300J**試験問題集を提供します。JPNTTest.com DP-300J試験問題集はもう更新されました。ここで**DP-300J**問題集のテストエンジンを手に入れます。最新版のアクセス、<https://www.jpntest.com/shiken/DP-300J-mondaishu> **503**問、**30%ディスカウント**、特別な割引コード: **JPNshiken**」