QUESTIONS & ANSWERS

Kill your exam at first Attempt

KILL EXAMS

DP-200 Dumps
DP-200 Braindumps
DP-200 Real Questions
DP-200 Practice Test
DP-200 dumps free

## Microsoft

# DP-200

*Implementing an Azure Data Solution*

Question #3 Section 19

Introductory Info Overview -

Current environment -
Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.
The majority of the company:€™s data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

| Applications | Tier | Replication | Notes |
|---|---|---|---|
| Internal Contoso | 1 | Yes | |
| Internal Contoso | 2 | SQL Data Sync | Data Sync between databases |
| Internal Partner | 3 | Yes | Replicate to Partner |
| External Contoso | 4,5,6 | Yes | |
| External Partner | 7,8 | No | Partner managed |
| Internal Distribution and Sales | 9 | Yes, once ingested at branches | Data ingested from Contoso branches |
| External Distribution and Sales | 10 | Yes, once ingested at Contoso main office | Data is ingested from multiple sources |

The company has a reporting infrastructure that ingests data from local databases and partner services. Partners services consists of distributors, wholesales, and retailers across the world. The company performs daily, weekly, and monthly reporting.

Requirements -
Tier 3 and Tier 6 through Tier 8 application must use database density on the same server and Elastic pools in a cost-effective manner.
Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.
A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of server going offline.
Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.
Tier 1 internal applications on the premium P2 tier
Tier 2 internal applications on the standard S4 tier
The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.
Tier 7 and Tier 8 partner access must be restricted to the database only.
In addition to default Azure backup behavior, Tier 4 and 5 databases must be on a backup strategy that performs a transaction log backup eve hour, a differential backup of databases every day and a full back up every week.
Back up strategies must be put in place for all other standalone Azure SQL Databases using Azure SQL-provided backup storage and capabilities.

Databases -
Contoso requires their data estate to be designed and implemented in the Azure Cloud. Moving to the cloud must not inhibit access to or availability of data.
Databases:
Tier 1 Database must implement data masking using the following masking logic:

| Data type | Masking requirement |
|---|---|
| A | Mask 4 or less string data type characters |
| B | Mask first letter and domain |
| C | Mask everything except characters at the beginning and end |

Tier 2 databases must sync between branches and cloud databases and in the event of conflicts must be set up for conflicts to be won by on-premises databases.
Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.
Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.
A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of a server going offline.
Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.
Tier 1 internal applications on the premium P2 tier
Tier 2 internal applications on the standard S4 tier

Reporting -

Security and monitoring -

Security -
A method of managing multiple databases in the cloud at the same time is must be implemented to streamlining data management and limiting management access to only those requiring access.

Monitoring -
Monitoring must be set up on every database. Contoso and partners must receive performance reports as part of contractual agreements.
Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.
The Azure SQL Data Warehouse cache must be monitored when the database is being used. A dashboard monitoring key performance indicators (KPIs) indicated by traffic lights must be created and displayed based on the following metrics:

| Metric | Description |
|---|---|
| A | Low cache hit %, high cache usage % |
| B | Low cache hit %, low cache usage % |
| C | High cache hit %, high cache usage % |

Existing Data Protection and Security compliances require that all certificates and keys are internally managed in an on-premises storage.
You identify the following reporting requirements:
Azure Data Warehouse must be used to gather and query data from multiple internal and external databases
Azure Data Warehouse must be optimized to use data from a cache
Reporting data aggregated for external partners must be stored in Azure Storage and be made available during regular business hours in the connecting regions
Reporting strategies must be improved to real time or near real time reporting cadence to improve competitiveness and the general supply chain
Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company¤€™s main office
Tier 10 reporting data must be stored in Azure Blobs

Issues -
Team members identify the following issues:
Both internal and external client application run complex joins, equality searches and group-by clauses. Because some systems are managed externally, the queries will not be changed or optimized by Contoso
External partner organization data formats, types and schemas are controlled by the partner companies
▪
Internal and external database development staff resources are primarily SQL developers familiar with the Transact-SQL language.
Size and amount of data has led to applications and reporting solutions not performing are required speeds
Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
The company maintains several legacy client applications. Data for these applications remains isolated form other applications. This has led to hundreds of databases being provisioned on a per application basis Question Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You need to configure data encryption for external applications.
Solution:
1. Access the Always Encrypted Wizard in SQL Server Management Studio
2. Select the column to be encrypted
3. Set the encryption type to Deterministic
4. Configure the master key to use the Azure Key Vault
5. Validate configuration results and deploy the solution
Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A
We use the Azure Key Vault, not the Windows Certificate Store, to store the master key.
Note: The Master Key Configuration page is where you set up your CMK (Column Master Key) and select the key store provider where the CMK will be stored.
Currently, you can store a CMK in the Windows certificate store, Azure Key Vault, or a hardware security module (HSM).

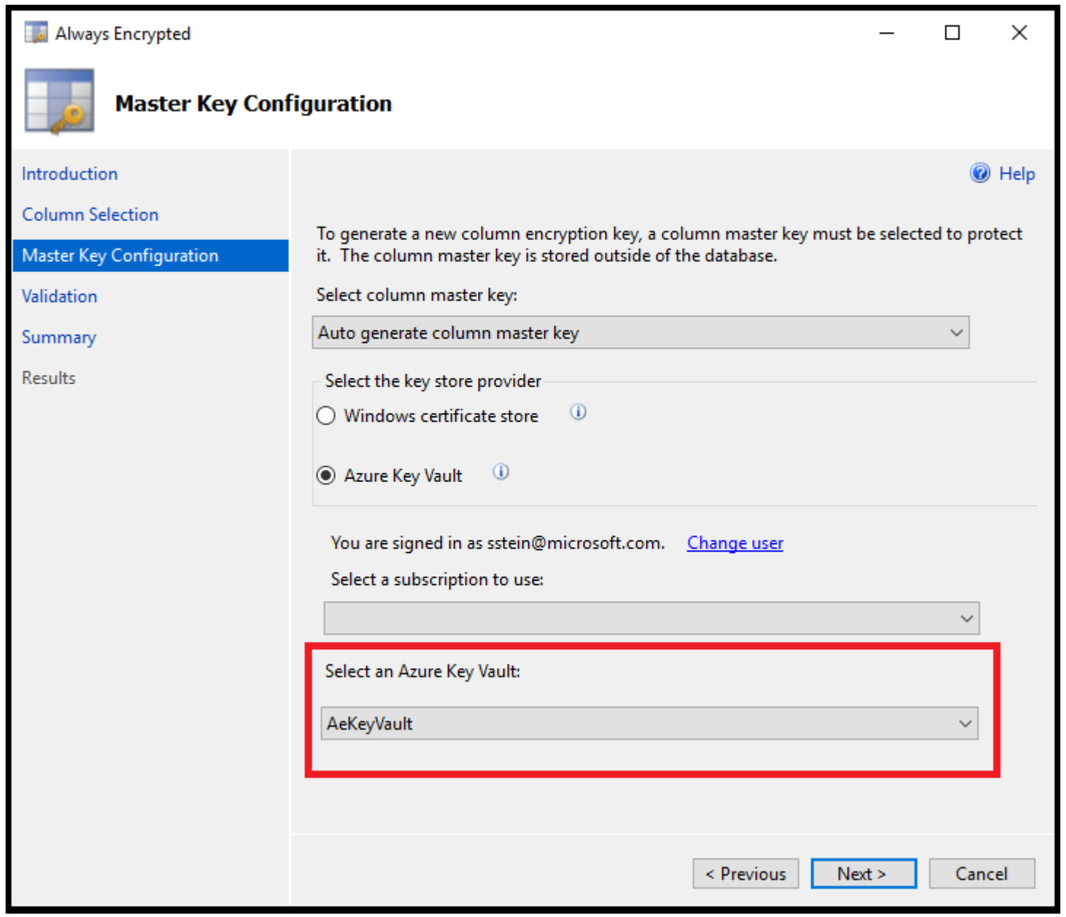


References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-always-encrypted-azure-key-vault

Question #4 Section 19

Introductory Info Overview -

Current environment -
Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.
The majority of the company‚€™s data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

| Applications | Tier | Replication | Notes |
|---|---|---|---|
| Internal Contoso | 1 | Yes | |
| Internal Contoso | 2 | SQL Data Sync | Data Sync between databases |
| Internal Partner | 3 | Yes | Replicate to Partner |
| External Contoso | 4,5,6 | Yes | |
| External Partner | 7,8 | No | Partner managed |
| Internal Distribution and Sales | 9 | Yes, once ingested at branches | Data ingested from Contoso branches |
| External Distribution and Sales | 10 | Yes, once ingested at Contoso main office | Data is ingested from multiple sources |

The company has a reporting infrastructure that ingests data from local databases and partner services. Partners services consists of distributors, wholesales, and retailers across the world. The company performs daily, weekly, and monthly reporting.

Requirements -
Tier 3 and Tier 6 through Tier 8 application must use database density on the same server and Elastic pools in a cost-effective manner.
Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.
A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of server going offline.
Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.
Tier 1 internal applications on the premium P2 tier
Tier 2 internal applications on the standard S4 tier
The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.
Tier 7 and Tier 8 partner access must be restricted to the database only.
In addition to default Azure backup behavior, Tier 4 and 5 databases must be on a backup strategy that performs a transaction log backup eve hour, a differential backup of databases every day and a full back up every week.
Back up strategies must be put in place for all other standalone Azure SQL Databases using Azure SQL-provided backup storage and capabilities.

Databases -
Contoso requires their data estate to be designed and implemented in the Azure Cloud. Moving to the cloud must not inhibit access to or availability of data.
Databases:
Tier 1 Database must implement data masking using the following masking logic:

| Data type | Masking requirement |
|---|---|
| A | Mask 4 or less string data type characters |
| B | Mask first letter and domain |
| C | Mask everything except characters at the beginning and end |

Tier 2 databases must sync between branches and cloud databases and in the event of conflicts must be set up for conflicts to be won by on-premises databases.
Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.
Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.
A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of a server going offline.
Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.
Tier 1 internal applications on the premium P2 tier
Tier 2 internal applications on the standard S4 tier

Reporting -

Security and monitoring -

Security -
A method of managing multiple databases in the cloud at the same time is must be implemented to streamlining data management and limiting management access to only those requiring access.

Monitoring -
Monitoring must be set up on every database. Contoso and partners must receive performance reports as part of contractual agreements.
Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.
The Azure SQL Data Warehouse cache must be monitored when the database is being used. A dashboard monitoring key performance indicators (KPIs) indicated by traffic lights must be created and displayed based on the following metrics:

| Metric | Description |
|---|---|
| A | Low cache hit %, high cache usage % |
| B | Low cache hit %, low cache usage % |
| C | High cache hit %, high cache usage % |

Existing Data Protection and Security compliances require that all certificates and keys are internally managed in an on-premises storage.
You identify the following reporting requirements:
Azure Data Warehouse must be used to gather and query data from multiple internal and external databases
Azure Data Warehouse must be optimized to use data from a cache
Reporting data aggregated for external partners must be stored in Azure Storage and be made available during regular business hours in the connecting regions
Reporting strategies must be improved to real time or near real time reporting cadence to improve competitiveness and the general supply chain
Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company¤€™s main office
Tier 10 reporting data must be stored in Azure Blobs

Issues -
Team members identify the following issues:
Both internal and external client application run complex joins, equality searches and group-by clauses. Because some systems are managed externally, the queries will not be changed or optimized by Contoso
External partner organization data formats, types and schemas are controlled by the partner companies
▪
Internal and external database development staff resources are primarily SQL developers familiar with the Transact-SQL language.
Size and amount of data has led to applications and reporting solutions not performing are required speeds
Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
The company maintains several legacy client applications. Data for these applications remains isolated form other applications. This has led to hundreds of databases being provisioned on a per application basis Question HOTSPOT -
You need to mask tier 1 data. Which functions should you use? To answer, select the appropriate option in the answer area.
NOTE: Each correct selection is worth one point.
Hot Area:

**Answer Area**

| Data type | Masking function |
| --- | --- |
| A | custom text / default / email / random number |
| B | custom text / default / email / random number |
| C | custom text / default / email / random number |

**Answer Area**

| Data type | Masking function |
|---|---|
| A | custom text / default / email / random number |
| B | custom text / default / email / random number |
| C | custom text / default / email / random number |

**Answer:**

A: Default -
Full masking according to the data types of the designated fields.
For string data types, use XXXX or fewer Xs if the size of the field is less than 4 characters (char, nchar, varchar, nvarchar, text, ntext).

B: email -

C: Custom text -
Custom StringMasking method which exposes the first and last letters and adds a custom padding string in the middle. prefix,[padding],suffix
Tier 1 Database must implement data masking using the following masking logic:

| Data type | Masking requirement |
|---|---|
| A | Mask 4 or less string data type characters |
| B | Mask first letter and domain |
| C | Mask everything except characters at the beginning and end |

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

Question #5 Section 19

Introductory Info Overview -

Current environment -
Contoso relies on an extensive partner network for marketing, sales, and distribution. Contoso uses external companies that manufacture everything from the actual pharmaceutical to the packaging.
The majority of the company:€™s data reside in Microsoft SQL Server database. Application databases fall into one of the following tiers:

| Applications | Tier | Replication | Notes |
|---|---|---|---|
| Internal Contoso | 1 | Yes | |
| Internal Contoso | 2 | SQL Data Sync | Data Sync between databases |
| Internal Partner | 3 | Yes | Replicate to Partner |
| External Contoso | 4,5,6 | Yes | |
| External Partner | 7,8 | No | Partner managed |
| Internal Distribution and Sales | 9 | Yes, once ingested at branches | Data ingested from Contoso branches |
| External Distribution and Sales | 10 | Yes, once ingested at Contoso main office | Data is ingested from multiple sources |

The company has a reporting infrastructure that ingests data from local databases and partner services. Partners services consists of distributors, wholesales, and retailers across the world. The company performs daily, weekly, and monthly reporting.

Requirements -
Tier 3 and Tier 6 through Tier 8 application must use database density on the same server and Elastic pools in a cost-effective manner.
Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.
A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of server going offline.
Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.
Tier 1 internal applications on the premium P2 tier
Tier 2 internal applications on the standard S4 tier
The solution must support migrating databases that support external and internal application to Azure SQL Database. The migrated databases will be supported by Azure Data Factory pipelines for the continued movement, migration and updating of data both in the cloud and from local core business systems and repositories.
Tier 7 and Tier 8 partner access must be restricted to the database only.
In addition to default Azure backup behavior, Tier 4 and 5 databases must be on a backup strategy that performs a transaction log backup eve hour, a differential backup of databases every day and a full back up every week.
Back up strategies must be put in place for all other standalone Azure SQL Databases using Azure SQL-provided backup storage and capabilities.

Databases -
Contoso requires their data estate to be designed and implemented in the Azure Cloud. Moving to the cloud must not inhibit access to or availability of data.
Databases:
Tier 1 Database must implement data masking using the following masking logic:

| Data type | Masking requirement |
|---|---|
| A | Mask 4 or less string data type characters |
| B | Mask first letter and domain |
| C | Mask everything except characters at the beginning and end |

Tier 2 databases must sync between branches and cloud databases and in the event of conflicts must be set up for conflicts to be won by on-premises databases.
Tier 3 and Tier 6 through Tier 8 applications must use database density on the same server and Elastic pools in a cost-effective manner.
Applications must still have access to data from both internal and external applications keeping the data encrypted and secure at rest and in transit.
A disaster recovery strategy must be implemented for Tier 3 and Tier 6 through 8 allowing for failover in the case of a server going offline.
Selected internal applications must have the data hosted in single Microsoft Azure SQL Databases.
Tier 1 internal applications on the premium P2 tier
Tier 2 internal applications on the standard S4 tier

Reporting -

Security and monitoring -

Security -
A method of managing multiple databases in the cloud at the same time is must be implemented to streamlining data management and limiting management access to only those requiring access.

Monitoring -
Monitoring must be set up on every database. Contoso and partners must receive performance reports as part of contractual agreements.
Tiers 6 through 8 must have unexpected resource storage usage immediately reported to data engineers.
The Azure SQL Data Warehouse cache must be monitored when the database is being used. A dashboard monitoring key performance indicators (KPIs) indicated by traffic lights must be created and displayed based on the following metrics:

| Metric | Description |
|---|---|
| A | Low cache hit %, high cache usage % |
| B | Low cache hit %, low cache usage % |
| C | High cache hit %, high cache usage % |

Existing Data Protection and Security compliances require that all certificates and keys are internally managed in an on-premises storage.
You identify the following reporting requirements:
Azure Data Warehouse must be used to gather and query data from multiple internal and external databases
Azure Data Warehouse must be optimized to use data from a cache
Reporting data aggregated for external partners must be stored in Azure Storage and be made available during regular business hours in the connecting regions
Reporting strategies must be improved to real time or near real time reporting cadence to improve competitiveness and the general supply chain
Tier 9 reporting must be moved to Event Hubs, queried, and persisted in the same Azure region as the company¤€™s main office
Tier 10 reporting data must be stored in Azure Blobs

Issues -
Team members identify the following issues:
Both internal and external client application run complex joins, equality searches and group-by clauses. Because some systems are managed externally, the queries will not be changed or optimized by Contoso
External partner organization data formats, types and schemas are controlled by the partner companies
▪
Internal and external database development staff resources are primarily SQL developers familiar with the Transact-SQL language.
Size and amount of data has led to applications and reporting solutions not performing are required speeds
Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
The company maintains several legacy client applications. Data for these applications remains isolated form other applications. This has led to hundreds of databases being provisioned on a per application basis Question DRAG DROP -
You need to set up access to Azure SQL Database for Tier 7 and Tier 8 partners.
Which three actions should you perform in sequence? To answer, move the appropriate three actions from the list of actions to the answer area and arrange them in the correct order.
Select and Place:

**Actions**

- Connect to the Database and use Azure PowerShell to create a database firewall rule
- Set the Allow Azure Services to Access Server to Disabled
- In ther Azure portal, create a database firewall rule
- In the Azure portal, create a server firewall rule
- Connect to the database and use Transact-SQL to create a database firewall rule
- Set the Allow Azure Services to Access Server setting to Enabled

**Answer Area**

**Actions**

- Connect to the Database and use Azure PowerShell to create a database firewall rule
- Set the Allow Azure Services to Access Server to Disabled
- In ther Azure portal, create a database firewall rule
- In the Azure portal, create a server firewall rule
- Connect to the database and use Transact-SQL to create a database firewall rule
- Set the Allow Azure Services to Access Server setting to Enabled

**Answer Area**

- Set the Allow Azure Services to Access Server to Disabled
- In the Azure portal, create a server firewall rule
- Connect to the database and use Transact-SQL to create a database firewall rule

**Answer:**
Tier 7 and 8 data access is constrained to single endpoints managed by partners for access
Step 1: Set the Allow Azure Services to Access Server setting to Disabled
Set Allow access to Azure services to OFF for the most secure configuration.
By default, access through the SQL Database firewall is enabled for all Azure services, under Allow access to Azure services. Choose OFF to disable access for all Azure services.
Note: The firewall pane has an ON/OFF button that is labeled Allow access to Azure services. The ON setting allows communications from all Azure IP addresses and all Azure subnets. These Azure IPs or subnets might not be owned by you. This ON setting is probably more open than you want your SQL Database to be.
The virtual network rule feature offers much finer granular control.
Step 2: In the Azure portal, create a server firewall rule
Set up SQL Database server firewall rules
Server-level IP firewall rules apply to all databases within the same SQL Database server.
To set up a server-level firewall rule:
1. In Azure portal, select SQL databases from the left-hand menu, and select your database on the SQL databases page.
2. On the Overview page, select Set server firewall. The Firewall settings page for the database server opens.
Step 3: Connect to the database and use Transact-SQL to create a database firewall rule
Database-level firewall rules can only be configured using Transact-SQL (T-SQL) statements, and only after you've configured a server-level firewall rule.
To setup a database-level firewall rule:
1. Connect to the database, for example using SQL Server Management Studio.
2. In Object Explorer, right-click the database and select New Query.
3. In the query window, add this statement and modify the IP address to your public IP address:
EXECUTE sp_set_database_firewall_rule N'Example DB Rule','0.0.0.4','0.0.0.4';
4. On the toolbar, select Execute to create the firewall rule.
References:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-security-tutorial
Manage data security