



Microsoft

AZ-305

**Designing Microsoft Azure Infrastructure Solutions
QUESTION & ANSWERS**

QUESTION 1

A company needs a datastore created in Azure for an application. Below are the key requirements for the data store.

Ability to store JSON based items

Ability to use SQL like queries on the datastore

Ability to provide low latency access to data items

Which of the following would you consider as the data store?

- A. Azure BLOB storage
- B. Azure CosmosDB
- C. Azure HDInsight
- D. Azure Redis

Correct Answer: B

Explanation/Reference:

You can use CosmosDB to provide low latency access to data. You can use the SQL API to store JSON based objects. The Microsoft documentation mentions the following.

SQL query examples for Azure Cosmos DB

04/04/2019 • 46 minutes to read • Contributors     all

Azure Cosmos DB SQL API accounts support querying items using Structured Query Language (SQL) as a JSON query language. The design goals of the Azure Cosmos DB query language are to:

- Support SQL, one of the most familiar and popular query languages, instead of inventing a new query language. SQL provides a formal programming model for rich queries over JSON items.

Option A is incorrect since this is used for object level storage.

Option C is incorrect since this is used for open source analytics.

Option D is incorrect since this is used for storing data in a low memory cache.

For more information on how to use SQL queries, please visit the below URL-

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-sql-query>

QUESTION 2

You have to design a Data Engineering solution for your company. The company currently has an Azure subscription. They also have application data hosted in a database on a Microsoft SQL Server hosted in their on-premises data center server. They want to implement the following requirements

Transfer transactional data from the on-premises SQL server onto a data warehouse in Azure.

Data needs to be transferred every day in the night as a scheduled job

A managed Spark cluster needs to be in place for data engineers to perform analysis on the data stored in the SQL data warehouse. Here the data engineers should have the ability to develop notebooks in Scale, R and Python.

They also need to have a data lake store in place for the ingestion of data from multiple data sources
Which of the following would the use for hosting the data warehouse in Azure?

- A. Azure Data Factory
- B. Azure Databricks
- C. Azure Data Lake Gen2 Storage accounts
- D. Azure Synapse Analytics

Correct Answer: D

Explanation/Reference:

Here we can host a data warehouse in Azure with the help of the Azure Synapse Analytics service.
Option A is incorrect since this is used as data integration service in Azure
Option B is incorrect since this is used as a managed Data bricks service in Azure. This can be used as a data analytics platform
Option C is incorrect since this is used as a data lake store in Azure
For more information on Azure Synapse, one can go to the following URL
<https://docs.microsoft.com/en-us/azure/synapse-analytics/>

QUESTION 3

Your company currently has an application that is hosted on their on-premises environment. The application currently connects to two databases in the on-premises environment. The databases are named whizlabdb1 and whizlabdb2.
You have to move the databases onto Azure. The databases have to support server-side transactions across both of the databases.
You decide to deploy the databases to an Azure SQL database-managed instance.
Would this fulfill the requirement?

- A. Yes
- B. No

Correct Answer: A

Explanation/Reference:

When it comes to distributed transactions, this is supported for the Azure SQL Managed Instance. If there are multiple instances, you need to create a Server Trust Group and the instances need to be added to the Server Trust group.
For more information on distributed transactions in the cloud, one can visit the following URL
<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-transactions-overview>

QUESTION 4

Your company has an on-premises Hyper-V cluster that contains 20 virtual machines. Some of the virtual machines are based on Windows and some in Linux. You have to migrate the virtual machines onto Azure.

You have to recommend a solution that would be used to replicate the disks of the virtual machines to Azure. The solution needs to ensure that the virtual machines remain available when the migration of the disks is in progress.

You decide to create an Azure storage account and then run AzCopy

Would this fulfill the requirement?

- A. Yes
- B. No

Correct Answer: B

Explanation/Reference:

For this requirement, you should either use the Azure Migrate or Azure Site Recovery service

For more information on Azure Site Recovery and Azure Migrate, you can visit the below link

<https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-overview>

<https://docs.microsoft.com/en-us/azure/migrate/migrate-services-overview>

QUESTION 5

Case Study

An insurance company, HABInsurance, operates in three states and provides home, auto, and boat insurance. Besides the head office, HABInsurance has three regional offices.

Current environment

General

An insurance company, HABInsurance, operates in three states and provides home, auto, and boat insurance. Besides the head office, HABInsurance has three regional offices.

Technology assessment

The company has two Active Directory forests: main.habinsurance.com and region.habinsurance.com.

HABInsurance's primary internal system is Insurance Processing System (IPS). It is an ASP.Net/C# application running on IIS/Windows Servers hosted in a data center. IPS has three tiers: web, business logic API, and a datastore on a back end. The company uses Microsoft SQL Server and MongoDB for the backend. The system has two parts: Customer data and Insurance forms and documents.

Customer data is stored in Microsoft SQL Server and Insurance forms and documents — in MongoDB.

The company also has 10 TB of Human Resources (HR) data stored on NAS at the head office location.

Requirements

General

HABInsurance plans to migrate its workloads to Azure. They purchased an Azure subscription. Changes

During a transition period, HABInsurance wants to create a hybrid identity model along with a

Microsoft Office 365 deployment. The company intends to sync its AD forests to Azure AD and benefit from Azure AD administrative units functionality.

HABInsurance needs to migrate the current IPSCustomers SQL database to a new fully managed SQL database in Azure that would be budget-oriented, balanced with scalable compute and storage options. The management team expects the Azure database service to scale the database resources dynamically with minimal downtime. The technical team proposes implementing a DTU-based purchasing model for the new database.

HABInsurance wants to migrate Insurance forms and documents to Azure database service.

HABInsurance plans to move IPS first two tiers to Azure without any modifications. The technology team discusses the possibility of running IPS tiers on a set of virtual machines instances. The number of instances should be adjusted automatically based on the CPU utilization. An SLA of 99.95% must be guaranteed for the compute infrastructure.

The company needs to move HR data to Azure File shares.

In their new Azure ecosystem, HABInsurance plans to use internal and third-party applications. The company considers adding user consent for data access to the registered applications

Later, the technology team contemplates adding a customer self-service portal to IPS and deploying a new IPS to multi-region ASK. But the management team is worried about performance and availability of the multi-region AKS deployments during regional outages.

What two parameters would you recommend set up to ensure that the new IPSCustomers database will scale to meet the workload demands?

- A. Define the maximum of CPU cores
- B. Define the maximum resource limit per group of databases
- C. Define the maximum of Database Transaction Units
- D. Define the maximum of the allocated storage
- E. Define the maximum size for a database

Correct Answer: C,E

Explanation/Reference:

From the requirements statement:

"The management team expects the Azure database service to scale the database resources dynamically with minimal downtime. The technical team proposes implementing a DTU-based purchasing model for the new database."

Azure provides dynamic scalability for the Azure SQL Databases and Azure Managed Instances. Azure Managed instance service is based only on the vCore purchasing model. Azure SQL Database service offers two purchasing models: DTU-based and vCore-based. There are two deployment options for Azure SQL Database: Single database and Elastic pool — a collection of the single databases with shared resources.

The DTU-based model allows you to select two scalability parameters for a single Azure SQL Database: the maximum of Database Transaction Units (DTU) and the maximum database size. You can dynamically change these two parameters.

The DTU-based Azure SQL Database has three service tiers: Basic, Standard, and Premium.

DTU-based purchasing model

Basic (For less demanding workloads)

Standard (For workloads with typical performance requirements)

Premium (For IO-intensive workloads)

You can select vCore and DTU tiers from the Service tier dropdown (Number 2) on the SQL Database Configuration screen (Number 1). Then, using sliders, you set or change the SQL database resources: the number of DTUs (Number 3) and the database size (Number 4). The default values for an initial Standard (S0) tier are 10 DTUs and 2 GB for the database size.

Microsoft Azure

Dashboard > SQL databases > Create SQL Database >

Configure

Feedback

Service and compute tier

Select from the available tiers based on the needs of your workload. The vCore model provides a wide range of configuration controls and offers Hyperscale and Serverless to automatically scale your database based on your workload needs. Alternately, the DTU model provides set price/performance packages to choose from for easy configuration. [Learn more](#)

Service tier: Standard (For workloads with typical performance requirements) (2)

DTUs: 10 (S0) (3)

Data max size (GB): 2 (4)

Cost summary	
Cost per DTU (in USD)	1.50
DTUs selected	x 10
ESTIMATED COST / MONTH	15.00 USD

Options A and D are incorrect because you define the maximum CPU cores and the maximum allocated storage for the vCore-based model but not for the DTU-based.

Option B is incorrect because you need to define a resource limit per group of databases to dynamically scale Azure SQL Database Elastic pools but not for the DTU-based Azure SQL Database dynamic scaling.

For more information about Azure SQL Database scalability, please visit the below URLs:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/scale-resources>

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

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

QUESTION 6

A company has an on-premises file server cbflserver that runs Windows Server 2019. Windows Admin Center manages this server. The company owns an Azure subscription. You need to provide an Azure solution to prevent data loss if the file server fails.

You decide to create an Azure Recovery Services vault. You then decide to install the Azure Backup

agent and then schedule the backup.
Would this meet the requirement?

- A. Yes
- B. No

Correct Answer: A

Explanation/Reference:

Azure Recovery Services Vault is storage for backup data. The data can come from various Azure services, like VMs, Azure SQL database, Azure Backup Services, etc. To bring the files to Azure, you need to install the Microsoft Azure Recovery Service (MARS) or Azure Backup agent on a Windows file server machine. Then you create backup policies, schedule a backup, and define the backup files retention. Once this is in place, the backup agent transfers the files from cbflserver to Azure Recovery Services Vault.

For more information about Azure Backup services for on-premises files, please visit the following URLs:

<https://docs.microsoft.com/en-us/azure/backup/backup-architecture#architecture-direct-backup-of-on-premises-windows-server-machines-or-azure-vm-files-or-folders>

<https://docs.microsoft.com/en-us/azure/backup/backup-windows-with-mars-agent>

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-recovery-services-vault-overview>

QUESTION 7

A company is planning on deploying an application onto Azure. The application will be based on the .Net core programming language. The application would be hosted using Azure Web apps. Below is part of the various requirements for the application

Give the ability to correlate Azure resource usage and the performance data with the actual application configuration and performance data

Give the ability to visualize the relationships between application components

Give the ability to track requests and exceptions to specific lines of code from within the application

Give the ability to actually analyse how users return to an application and see how often they only select a particular drop-down value

Which of the following service would be best suited for fulfilling the requirement of

“Give the ability to correlate Azure resource usage and the performance data with the actual application configuration and performance data”

- A. Azure Application Insights
- B. Azure Service Map
- C. Azure Log Analytics
- D. Azure Activity Log

Correct Answer: C

Explanation/Reference:

You can send data about the application and resource usage to Azure Log Analytics. You can then build queries on the stored data.

For more information on Azure Log Analytics, please go ahead and visit the below URL <https://docs.microsoft.com/en-us/azure/azure-monitor/learn/tutorial-viewdata>

QUESTION 8

A company has an on-premises file server cbflserver that runs Windows Server 2019. Windows Admin Center manages this server. The company owns an Azure subscription. You need to provide an Azure solution to prevent data loss if the file server fails.

You decide to register Windows Admin Center in Azure and then configure Azure Backup.

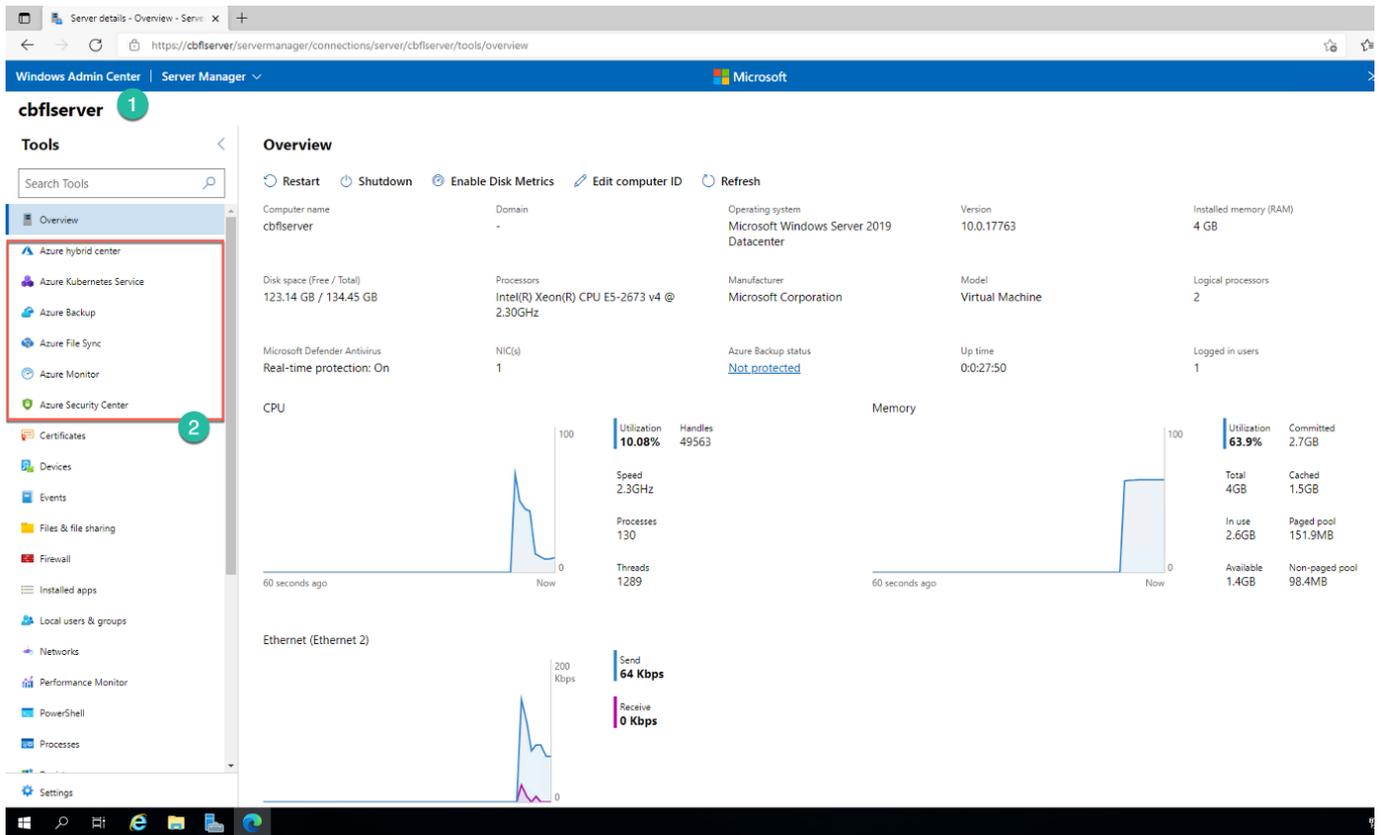
Would this meet the requirement?

- A. Yes
- B. No

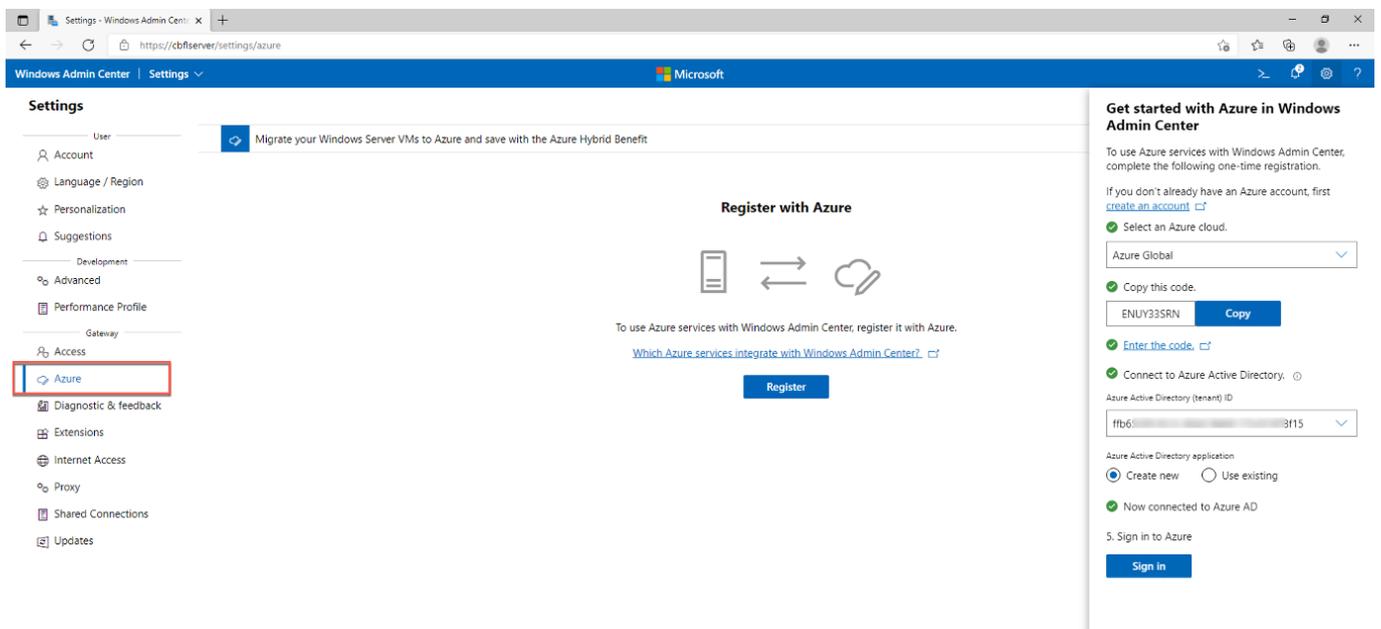
Correct Answer: A

Explanation/Reference:

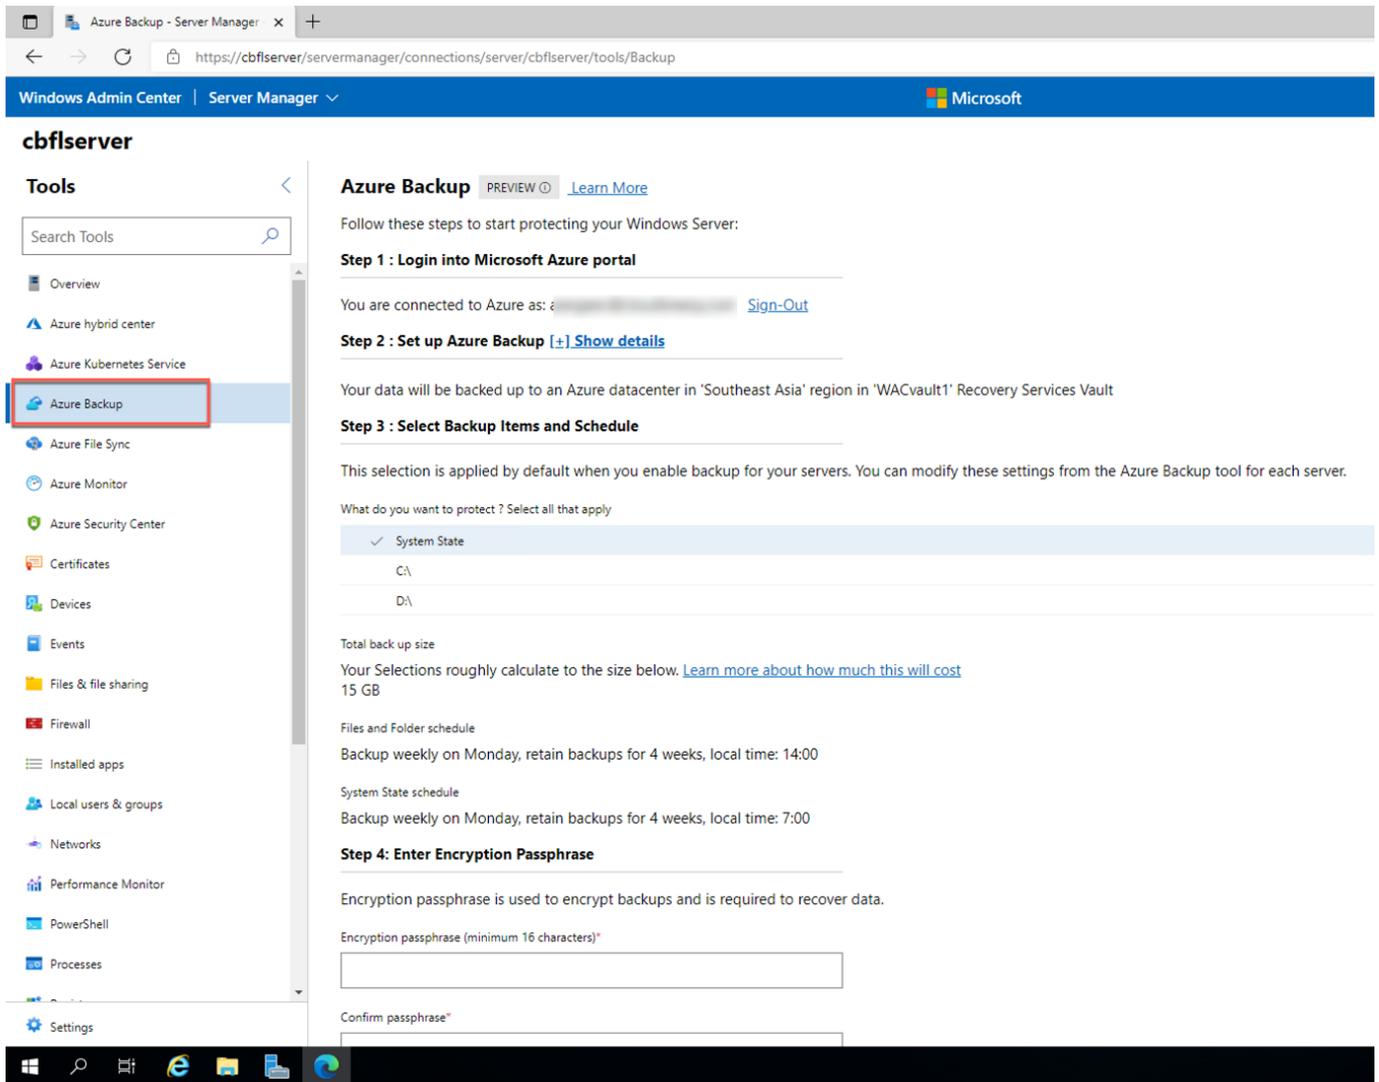
To bring the files to Azure, you can use Windows Admin Center with Azure Backup service. The Center is a locally deployed browser-based solution (Number 1). It manages on-premises Windows servers, clusters, etc. Windows Admin Center consists of two main components: the Gateway that manages the servers using PowerShell or WMI over WinRM and the Web server that listens to the HTTP requests and serves the UI to the browser. Azure integration is the built-in functionality of the Center (Number 2).



Before using Azure services, you need to connect and register the Center with Azure using a gateway.



Then you can set up and configure Azure Backup Service. You select the existent Recovery Services vault or create a new one during the setup, the vault region, and the Azure resource group. After you can add cbfserver files and directories to backup and create a backup schedule. And finally, enter an encryption phrase to encrypt backups. Once this is in place, the Center will launch the server files backup to Azure Recovery Services Vault.



For more information about Azure Backup services for on-premises files, please visit the following URLs:

<https://docs.microsoft.com/en-us/windows-server/manage/windows-admin-center/azure/azure-backup>

<https://docs.microsoft.com/en-us/azure/backup/backup-architecture#architecture-direct-backup-of-on-premises-windows-server-machines-or-azure-vm-files-or-folders>

<https://docs.microsoft.com/en-us/azure/backup/backup-windows-with-mars-agent>

QUESTION 9

Your company has an Azure Web App that runs via the Premium App Service Plan. A development team will be using the Azure Web App. You have to configure the Azure Web app so that it can fulfil the below requirements.

Provide the ability to switch the web app from the current version to a newer version

Provide developers with the ability to test newer versions of the application before the switch to the newer version occurs

Ensure that the application version can be rolled back

Minimize downtime

Which of the following can be used for this requirement?

- A. Create a new App Service Plan
- B. Make use of deployment slots
- C. Map a custom domain

D. Backup the Azure Web App

Correct Answer: B

Explanation/Reference:

With deployment slots, you can easily deploy newer versions of the applications onto the deployment slot.

Developers can test the newer version of the application in the deployment slot.

When the version needs to be promoted, the production slot can be switched with the deployment.

Rollbacks can also occur at any point in time.

Option A is incorrect since the Premium App Service plan has the option to use deployment slots. And using deployment slots is required for this question.

Option C is incorrect since there is no mention of using custom domain names.

Option D is incorrect since there is no mention of needing to back up the Azure Web App

For more information on deployment slots, one can visit the following URL

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots>

QUESTION 10

You have to deploy an Azure SQL database named udemydb for your company. The databases must meet the following security requirements

When IT help desk supervisors query a database table named customers, they must be able to see the full number of each credit card

When IT help desk operators query a database table named customers, they must only see the last four digits of each credit card number

A column named Credit Card rating in the customers table must never appear in plain text in the database system. Only client applications must be able to decrypt the information that is stored in this column

Which of the following can be implemented for the Credit Card rating column security requirement?

- A. Always Encrypted
- B. Azure Advanced Threat Protection
- C. Transparent Data Encryption
- D. Dynamic Data Masking

Correct Answer: A

Explanation/Reference:

You can use the Always Encrypted feature to ensure that data in the database tables are encrypted at rest. This will ensure that the data values never appear in plain text in the database system. Also only client applications that have the required encryption key will be able to decrypt the data

Option B is incorrect since this is used to check for any sort of threats to the underlying database

Option C is incorrect since this is used to encrypt data at rest

Option D is incorrect since this is used to mask data in the database tables

For more information on the Always Encrypted, one can go to the following URL

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine?view=sql-server-ver15>