QUESTION 23
You have 100 servers that run Windows Server 2012 R2 and host Microsoft SQL Server 2014 instances. The instances host databases that have the following characteristics:

- The largest database is currently 3 TB. None of the databases will ever exceed 4 TB.
- Stored procedures are implemented by using CLR.

You plan to move all the data from SQL Server to Azure.

You need to recommend an Azure service to host the databases. The solution must meet the following requirements:

- Whenever possible, minimize management overhead for the migrated databases.
- Minimize the number of database changes required to facilitate the migration.
- Ensure that users can authenticate by using their Active Directory credentials.

What should you include in the recommendation?

A. Azure SQL Database elastic pools
B. Azure SQL Database Managed Instance
C. Azure SQL Database single databases
D. SQL Server 2016 on Azure virtual machines

Correct Answer: B
Explanation

Explanation/Reference:
Reference:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-managed-instance

QUESTION 24
You are designing an order processing system in Azure that will contain the Azure resources shown in the following table.

Name	Type	Purpose
App1	Web app	Processes customer orders
Function1	Function	Check product availability at vendor 1
Function2	Function	Check product availability at vendor 2
storage1	Storage account	Stores order processing logs

The order processing system will have the following transaction flow:

- A customer will place an order by using App1.
- When the order is received, App1 will generate a message to check for product availability at vendor 1 and vendor 2.
- An integration component will process the message, and then trigger either Function1 or Function2 depending on the type of order. Once
- a vendor confirms the product availability, a status message for App1 will be generated by Function1 or Function2.
- All the steps of the transaction will be logged to storage1.

Which type of resource should you recommend for the integration component?

A. an Azure Data Factory pipeline
B. an Azure Service Bus queue
C. an Azure Event Grid domain
D. an Azure Event Hubs capture

Correct Answer: A
Explanation

Explanation/Reference:
Explanation:
A data factory can have one or more pipelines. A pipeline is a logical grouping of activities that together perform a

task. The activities in a pipeline define actions to perform on your data.

Data Factory has three groupings of activities: data movement activities, data transformation activities, and control activities.

Azure Functions is now integrated with Azure Data Factory, allowing you to run an Azure function as a step in your data factory

pipelines. Reference:
https://docs.microsoft.com/en-us/azure/data-factory/concepts-pipelines-activities

QUESTION 25
HOTSPOT

You have an existing implementation of Microsoft SQL Server Integration Services (SSIS) packages stored in an SSISDB catalog on your on-premises network. The on-premises network does not have hybrid connectivity to Azure by using Site-to-Site VPN or ExpressRoute.

You want to migrate the packages to Azure Data Factory.

You need to recommend a solution that facilitates the migration while minimizing changes to the existing packages. The solution must minimize

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Store the SSISDB catalog by using:
- Azure SQL Database
- Azure Synapse Analytics
- SQL Server on an Azure virtual machine
- SQL Server on an on-premises computer

Implement a runtime engine for package execution by using:
- Self-hosted integration runtime only
- Azure-SQL Server Integration Services Integration Runtime (IR) only
- Azure-SQL Server Integration Services Integration Runtime and self-hosted integration runtime

Correct Answer:

Answer Area

Store the SSISDB catalog by using:
- Azure SQL Database
- Azure Synapse Analytics
- SQL Server on an Azure virtual machine
- SQL Server on an on-premises computer

Implement a runtime engine for package execution by using:
- Self-hosted integration runtime only
- Azure-SQL Server Integration Services Integration Runtime (IR) only
- Azure-SQL Server Integration Services Integration Runtime and self-hosted integration runtime

Explanation

Explanation/Reference:
Explanation:

Box 1: Azure SQL database
You can't create the SSISDB Catalog database on Azure SQL Database at this time independently of creating the Azure-SSIS Integration Runtime in Azure Data Factory. The Azure-SSIS IR is the runtime environment that runs SSIS packages on Azure.

Box 2: Azure-SQL Server Integration Service Integration Runtime and self-hosted integration runtime

The Integration Runtime (IR) is the compute infrastructure used by Azure Data Factory to provide data integration capabilities across different network environments. Azure-SSIS Integration Runtime (IR) in Azure Data Factory (ADF) supports running SSIS packages.

Self-hosted integration runtime can be used for data movement in this

scenario. Reference:
https://docs.microsoft.com/en-us/azure/data-factory/create-azure-integration-runtime

https://docs.microsoft.com/en-us/sql/integration-services/lift-shift/ssis-azure-connect-to-catalog-database

QUESTION 26
You have 70 TB of files on your on-premises file server.

You need to recommend solution for importing data to Azure. The solution must minimize

cost. What Azure service should you recommend?

A. Azure StorSimple
B. Azure Batch
C. Azure Data Box
D. Azure Stack

Correct Answer: C
Explanation

Explanation/Reference:
Explanation:
Microsoft has engineered an extremely powerful solution that helps customers get their data to the Azure public cloud in a cost-effective, secure, and efficient manner with powerful Azure and machine learning at play. The solution is called Data Box.
Data Box and is in general availability status. It is a rugged device that allows organizations to have 100 TB of capacity on which to copy their data and then send it to be transferred to Azure.

Incorrect Answers:
A: StoreSimple would not be able to handle 70 TB of data.

Reference:
https://www.vembu.com/blog/what-is-microsoft-azure-data-box-disk-edge-heavy-gateway-overview/

QUESTION 27
You have an Azure subscription that contains 100 virtual machines.

You plan to design a data protection strategy to encrypt the virtual disks.

You need to recommend a solution to encrypt the disks by using Azure Disk Encryption. The solution must provide the ability to encrypt operating system disks and data disks.

What should you include in the recommendation?

A. a certificate
B. a key
C. a passphrase
D. a secret

Correct Answer: B
Explanation

Explanation/Reference:
Explanation:
For enhanced virtual machine (VM) security and compliance, virtual disks in Azure can be encrypted. Disks are encrypted by using cryptographic keys that are secured in an Azure Key Vault. You control these cryptographic keys and can audit their use.

Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/encrypt-disks

QUESTION 28
You have data files in Azure Blob storage.

You plan to transform the files and move them to Azure Data Lake

Storage. You need to transform the data by using mapping data flow.

Which Azure service should you use?

A. Azure Data Box Gateway
B. Azure Storage Sync
C. Azure Data Factory
D. Azure Databricks

Correct Answer: C
Explanation

Explanation/Reference:
Explanation:
You can use Copy Activity in Azure Data Factory to copy data from and to Azure Data Lake Storage Gen2, and use Data Flow to transform data in Azure Data Lake Storage Gen2.

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-data-lake-storage

QUESTION 29
You have an Azure virtual machine named VM1 that runs Windows Server 2019 and contains 500 GB of data files.

You are designing a solution that will use Azure Data Factory to transform the data files, and then load the files to Azure Data Lake

Storage. What should you deploy on VM1 to support the design?

A. the Azure Pipelines agent
B. the Azure File Sync agent
C. the On-premises data gateway
D. the self-hosted integration runtime in Azure

Correct Answer: D
Explanation

Explanation/Reference:
Explanation:
The integration runtime (IR) is the compute infrastructure that Azure Data Factory uses to provide data-integration capabilities across different network environments. For details about IR, see Integration runtime overview.

A self-hosted integration runtime can run copy activities between a cloud data store and a data store in a private network. It also can dispatch transform activities against compute resources in an on-premises network or an Azure virtual network. The installation of a self-hosted integration runtime needs an on-premises machine or a virtual machine inside a private network.

Reference:
https://docs.microsoft.com/en-us/azure/data-factory/create-self-hosted-integration-runtime

QUESTION 30
HOTSPOT

Your company is designing a multi-tenant application that will use elastic pools and Azure SQL databases. The application will be used by 30 customers.

You need to design a storage solution for the application. The solution must meet the following requirements:

- Operational costs must be minimized.
- All customers must have their own database.
- The customer databases will be in one of the following three Azure regions: East US, North Europe, or South Africa North.

What is the minimum number of elastic pools and Azure SQL Database servers required? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

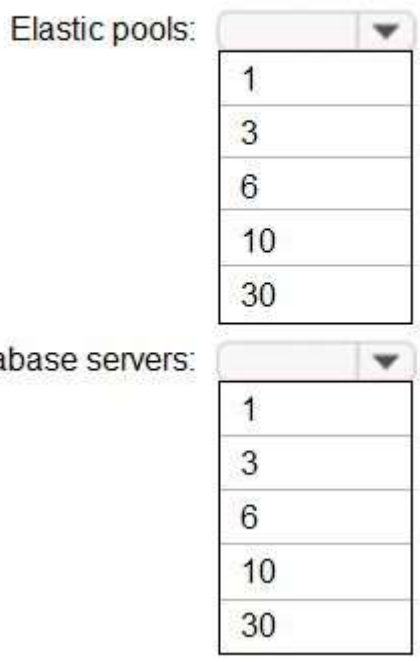

Correct Answer:

Answer Area

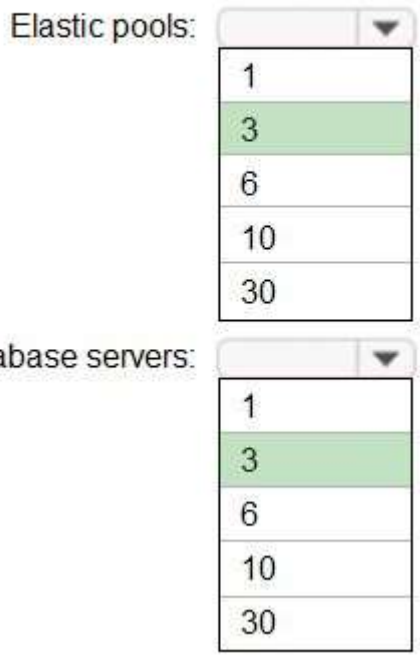

Explanation

Explanation/Reference:
Explanation:

Box 1: 3
The server, its pools & databases must be in the same Azure region under the same subscription.

Box 2: 3
A server can have up to 5000 databases associated to it.

Reference:
https://vincentlauzon.com/2016/12/18/azure-sql-elastic-pool-overview/

QUESTION 31
Your company purchases an app named App1.

You plan to run App1 on seven Azure virtual machines in an Availability Set. The number of fault domains is set to 3. The number of update domains is set to

20. You need to identify how many App1 instances will remain available during a period of planned maintenance.

How many App1 instances should you identify?

A. 1
B. 2
C. 6
D. 7

Correct Answer: C
Explanation

Explanation/Reference:
Explanation:
Only one update domain is rebooted at a time. Here there are 7 update domain with one VM each (and 13 update domain with no VM).

Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/manage-availability

QUESTION 32
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You have an Azure Storage v2 account named

storage1. You plan to archive data to storage1.

You need to ensure that the archived data cannot be deleted for five years. The solution must prevent administrators from deleting the

data. Solution: You create an Azure Blob storage container, and you configure a legal hold access policy.

Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Explanation

Explanation/Reference:
Explanation:
Use an Azure Blob storage container, but use a time-based retention policy instead of a legal hold.

Note:
Immutable storage for Azure Blob storage enables users to store business-critical data objects in a WORM (Write Once, Read Many) state. This state makes the data non-erasable and non-modifiable for a user-specified interval. For the duration of the retention interval, blobs can be created and read, but cannot be modified or deleted. Immutable storage is available for general-purpose v2 and Blob storage accounts in all Azure regions.

Note: Set retention policies and legal holds

1. Create a new container or select an existing container to store the blobs that need to be kept in the immutable state. The container must be in a general-purpose v2 or Blob storage account.

2. Select Access policy in the container settings. Then select Add policy under Immutable blob storage.

Either

3a. To enable legal holds, select Add Policy. Select Legal hold from the drop-down

menu. Or
3b. To enable time-based retention, select Time-based retention from the drop-down menu.

4. Enter the retention interval in days (acceptable values are 1 to 146000 days).

Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-immutable-storage

https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-immutability-policies-manage

QUESTION 33
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You have an Azure Storage v2 account named

storage1. You plan to archive data to storage1.

You need to ensure that the archived data cannot be deleted for five years. The solution must prevent administrators from deleting the

data. Solution: You create a file share and snapshots.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B
Explanation

Explanation/Reference:
Explanation:
Instead you could create an Azure Blob storage container, and you configure a legal hold access policy.

Reference:
https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-immutable-storage