DUMPSBOSS.

Designing Microsoft Azure Infrastructure Solutions

Microsoft AZ-305

Version Demo

Total Demo Questions: 15

Total Premium Questions: 279

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, New Update	220
Topic 2, Case Study 1	2
Topic 3, Case Study 2	2
Topic 4, Case Study 3	3
Topic 5, Mixed Questions	52
Total	279

QUESTION NO: 1

You have an Azure Active Directory (Azure AD) tenant that syncs with an on-premises Active Directory domain.

Your company has a line-of-business (LOB) application that was developed internally.

You need to implement. SAML single sign-on (SSO) and enforce multi-factor authentication (MFA) when users attempt to access the application from an unknown location.

Which two features should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Azure AD enterprise applications

B. Azure AD Identity Protection

C. Azure Application Gateway

D. Conditional Access policies

E. Azure AD Privileged Identity Management (PIM)

ANSWER: A D

Explanation:

To achieve SAML single sign-on (SSO) for an internally developed LOB application, you must configure Azure AD enterprise applications to integrate the application with Azure Active Directory using SAML-based authentication. Additionally, to enforce multi-factor authentication (MFA) when users are accessing the application from unknown locations, you should use Conditional Access policies. Conditional Access allows you to set conditions under which users must perform MFA to gain access to applications. This ensures added security when the system detects potentially risky sign-in behavior, such as those from unknown locations. More information can be found in Azure AD Connect documentation and Conditional Access documentation.

QUESTION NO: 2

You have an Azure Functions microservice app named Appl that is hosted in the Consumption plan. App1 uses an Azure Queue Storage trigger.

You plan to migrate App1 to an Azure Kubernetes Service (AKS) cluster.

You need to prepare the AKS cluster to support Appl. The solution must meet the following requirements:

• Use the same scaling mechanism as the current deployment.

• Support kubenet and Azure Container Netwoking Interface (CNI) networking.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Configure the horizontal pod autoscaler.

B. Install Virtual Kubelet.

C. Configure the AKS cluster autoscaler.

D. Configure the virtual node add-on.
Install Kubemetes-based Event Driven Autoscaling (KEDA).

ANSWER: A C D

Explanation:

To migrate App1 to AKS while ensuring that it uses the same scaling mechanisms, two crucial configurations are needed: - **Horizontal Pod Autoscaler**: This automatically scales the number of pods in a replication controller, deployment, or replica set based on observed CPU utilization or with custom metrics. This aligns with the scalability behavior in Azure Functions when hosted in a Consumption plan, which automatically scales based on demand. - **AKS Cluster Autoscaler**: This scales the nodes in a cluster automatically based on resource needs. This is akin to auto-scaling mechanisms in Azure Functions when the demand increases or decreases. For more details, check the official documentation on [AKS Scaling](https://learn.microsoft.com/en-us/azure/aks/cluster-autoscaler) and [Horizontal Pod Autoscaler](https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/).

QUESTION NO: 3

You are designing a large Azure environment that will contain many subscriptions.

You plan to use Azure Policy as part of a governance solution.

To which three scopes can you assign Azure Policy definitions? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Azure Active Directory (Azure AD) administrative units

B. Azure Active Directory (Azure AD) tenants

C. subscriptions

D. compute resources

E. resource groups

F. management groups

ANSWER: B C E F

Explanation:

Azure Policy evaluates resources in Azure by comparing the properties of those resources to business rules. Once your business rules have been formed, the policy definition or initiative is assigned to any scope of resources that Azure supports, such as management groups, subscriptions, resource groups, or individual resources.

Reference:

https://docs.microsoft.com/en-us/azure/governance/policy/overview

QUESTION NO: 4

You have an on-premises application named App1 that uses an Oracle database.

You plan to use Azure Databricks to transform and load data from App1 to an Azure Synapse Analytics instance.

You need to ensure that the App1 data is available to Databricks.

Which two Azure services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Data Box Edge

B. Azure Data Lake Storage

C. Azure Data Factory

D. Azure Data Box Gateway

E. Azure Import/Export service

ANSWER: B C

Explanation:

To enable data movement from an on-premises Oracle database to Azure Synapse Analytics using Azure Databricks, we can leverage Azure Data Factory and Azure Data Lake Storage. Azure Data Factory is a cloud-based ETL service that orchestrates data copy from on-premises data sources to cloud services. Azure Data Lake Storage Gen2 is optimized to store large amounts of structured and unstructured data, which enables Azure Databricks to process this data efficiently. Refer to Azure Data Factory and Azure Data Lake Storage for more details.

QUESTION NO: 5 - (HOTSPOT)

HOTSPOT

You have an on-premises database that you plan to migrate to Azure.

You need to design the database architecture to meet the following requirements:

- Support scaling up and down.
- Support geo-redundant backups.
- Support a database of up to 75 TB.
- Be optimized for online transaction processing (OLTP).

What should you include in the design? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Service:

- Azure SQL Database
- Azure SQL Managed Instance
- Azure Synapse Analytics
- SQL Server on Azure Virtual Machines

Service tier:

- Basic
- Business Critical
- General Purpose
- Hyperscale
- Premium
- Standard

ANSWER:

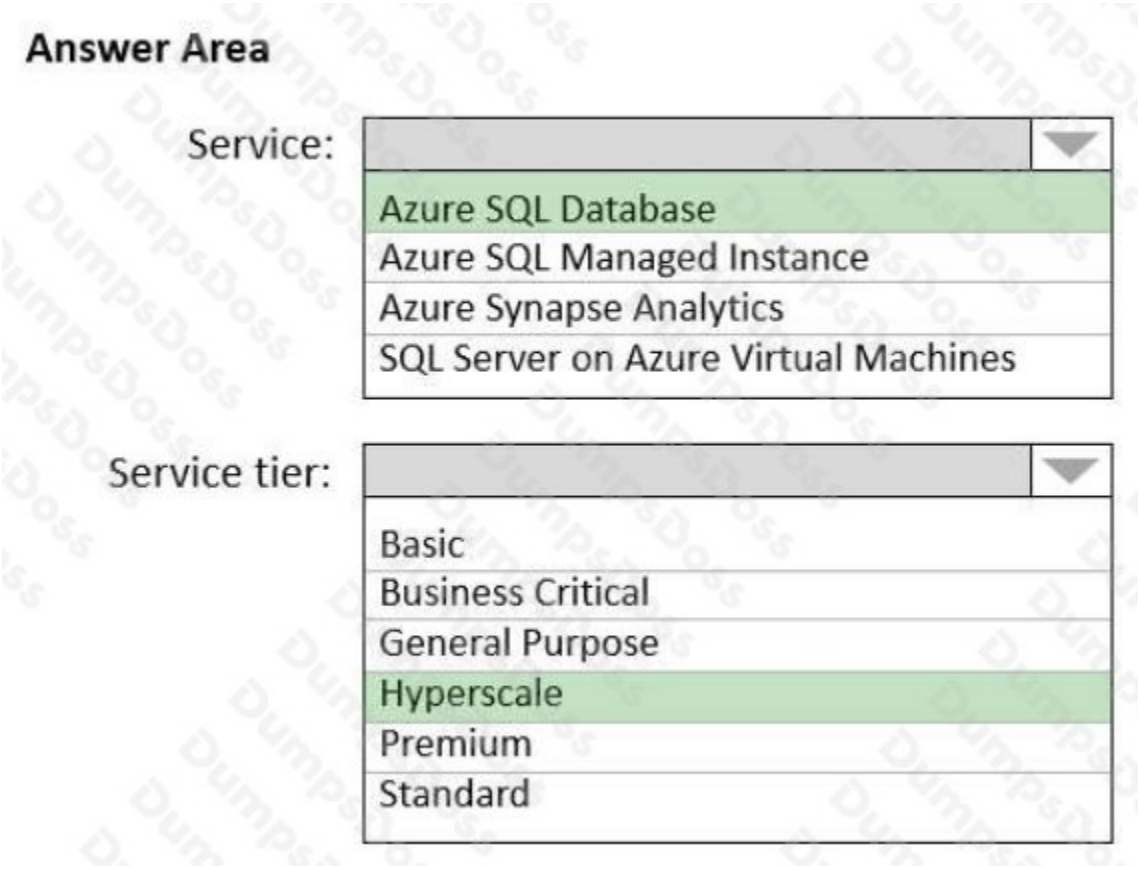

Explanation:

Box 1: Azure SQL Database Azure SQL Database:

Database size always depends on the underlying service tiers (e.g. Basic, Business Critical, Hyperscale). It supports databases of up to 100 TB with Hyperscale service tier model.

Active geo-replication is a feature that lets you to create a continuously synchronized readable secondary database for a primary database. The readable secondary database may be in the same Azure region as the primary, or, more commonly, in a different region. This kind of readable secondary databases are also known as geo-secondaries, or geo-replicas.

Azure SQL Database and SQL Managed Instance enable you to dynamically add more resources to your database with minimal downtime. Box 2: Hyperscale

Incorrect Answers:

- SQL Server on Azure VM: geo-replication not supported.
- Azure Synapse Analytics is not optimized for online transaction processing (OLTP).
- Azure SQL Managed Instance max database size is up to currently available instance size (depending on the number of vCores).

Max instance storage size (reserved) - 2 TB for 4 vCores

- 8 TB for 8 vCores - 16 TB for other sizes

Reference: https://docs.microsoft.com/en-us/azure/azure-sql/database/active-geo-replication-overview
https://medium.com/awesome-azure/azure-difference-between-azure-sql-database-and-sql-server-on-vm-comparison-azure-sql-vs-sql-server-vm-cf02578a1188

QUESTION NO: 6

You have an Azure subscription that contains an Azure SQL database.

You plan to use Azure reservations on the Azure SQL database.

To which resource type will the reservation discount be applied?

A. vCore compute

B. DTU compute

C. Storage

D. License

ANSWER: A

Explanation:

Quantity: The amount of compute resources being purchased within the capacity reservation. The quantity is a number of vCores in the selected Azure region and Performance tier that are being reserved and will get the billing discount. For example, if you run or plan to run multiple databases with the total compute capacity of Gen5 16 vCores in the East US region, then you would specify the quantity as 16 to maximize the benefit for all the databases.

Reference:

https://docs.microsoft.com/en-us/azure/azure-sql/database/reserved-capacity-overview

QUESTION NO: 7

You need to design a highly available Azure SQL database that meets the following requirements:

* Failover between replicas of the database must occur without any data loss.

* The database must remain available in the event of a zone outage.

* Costs must be minimized.

Which deployment option should you use?

A. Azure SQL Database Business Critical

B. Azure SQL Database Managed Instance Business Critical

C. Azure SQL Database Hyperscale

D. Azure SQL Database Standard

ANSWER: A

Explanation:

Azure SQL Database Business Critical provides high availability with automatic failover using the Always On availability groups technology without data loss, which satisfies the requirement of zero data loss in case of failover. It also supports zone redundancy, maintaining availability even in the event of a zone outage, ensuring high availability and disaster recovery. While there may be cheaper alternatives, the requirement for no data loss and resilience in a zone outage requires the features uniquely provided by the Business Critical service tier. For more information, please see the official documentation: Azure SQL Database Business Critical.

QUESTION NO: 8 - (DRAG DROP)

You need to design an architecture to capture the creation of users and the assignment of roles. The captured data must be stored in Azure Cosmos DB.

Which Azure services should you include in the design? To answer, drag the appropriate services to the correct targets. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

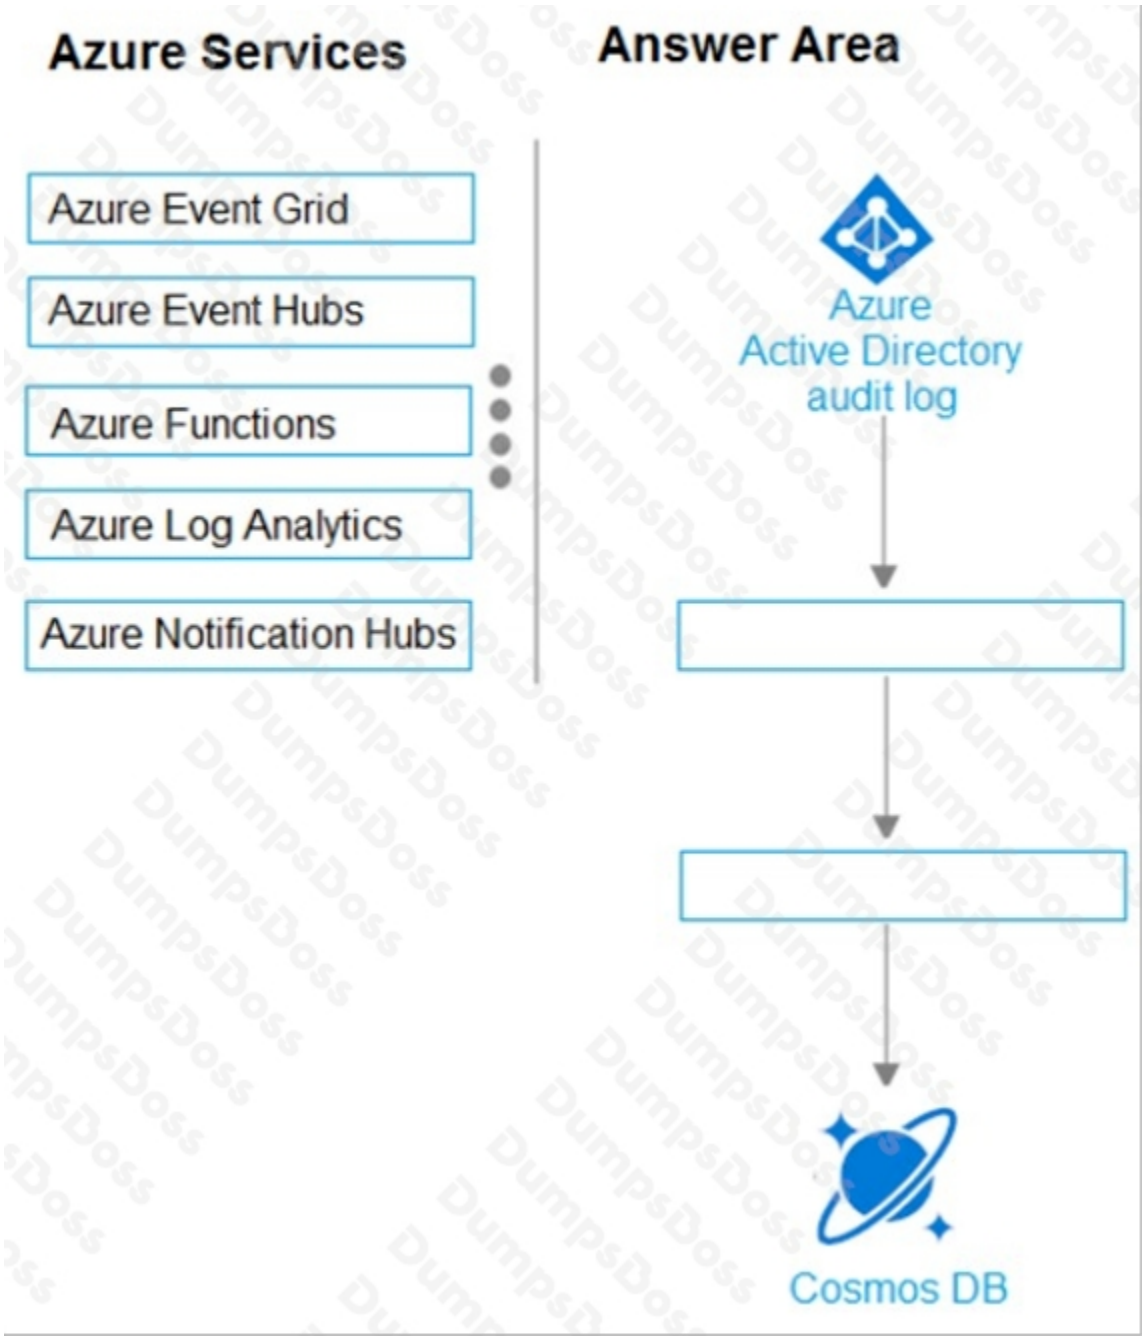
Azure Services
Azure Event Grid
Azure Event Hubs
Azure Functions
Azure Log Analytics
Azure Notification Hubs
Answer Area
Azure
Active Directory
audit log
Cosmos DB

ANSWER:

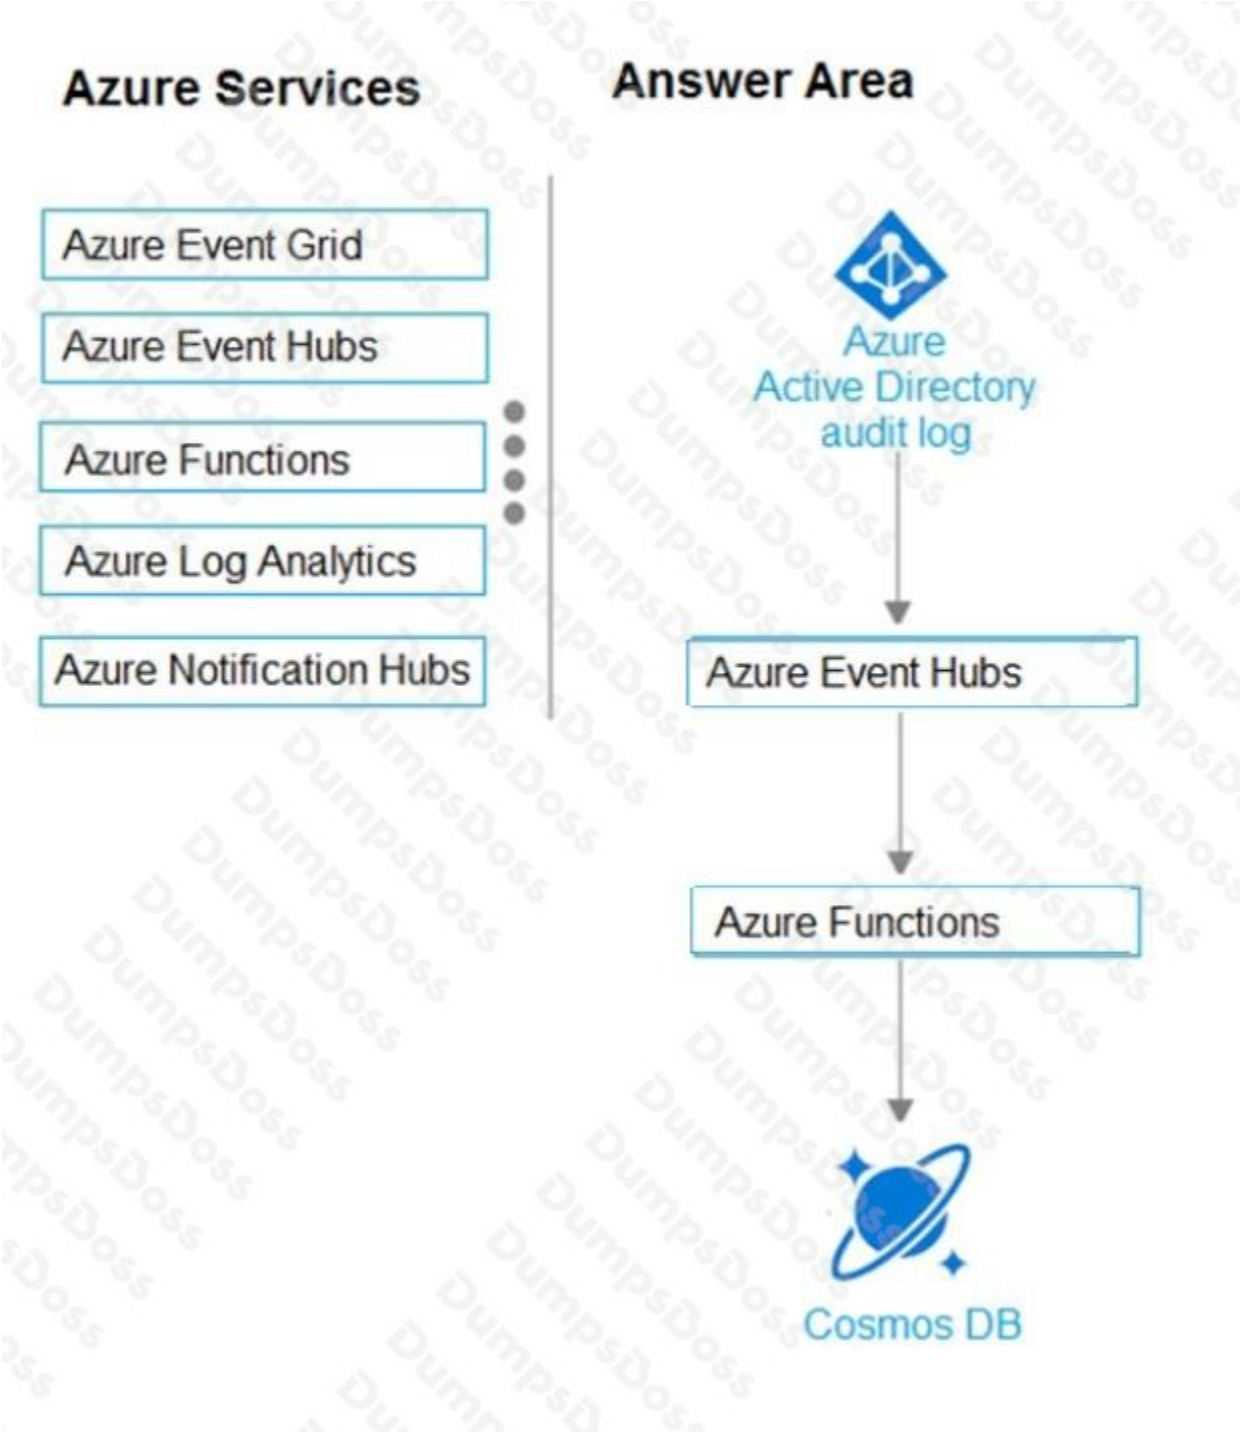

Explanation:

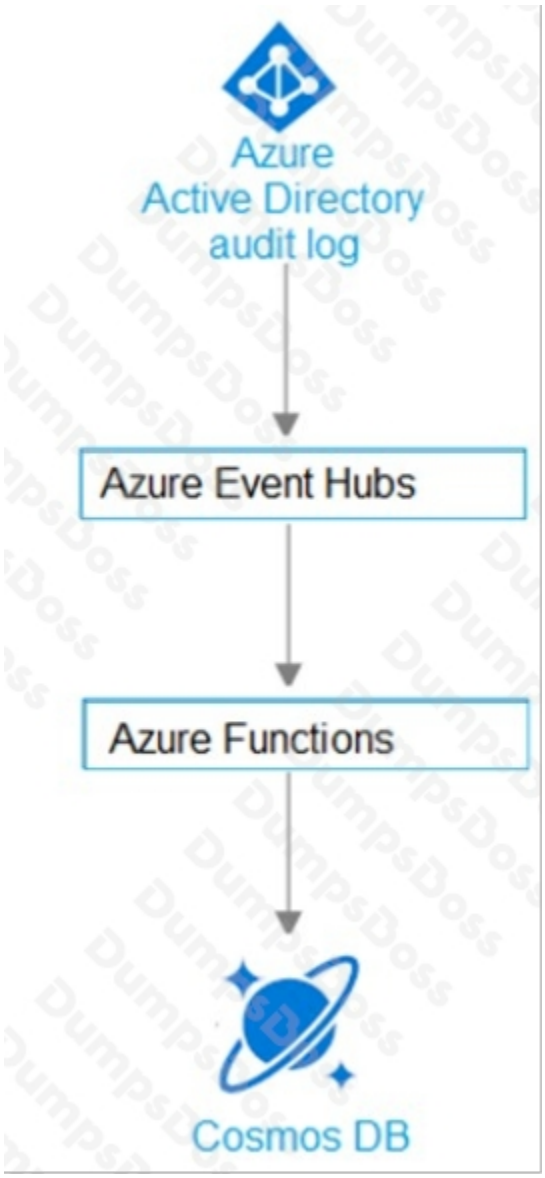

1. AAD audit log -> Event Hub (other two choices, LAW, storage, but not available in this question)

https://docs.microsoft.com/en-us/azure/active-directory/reports-monitoring/tutorial-azure-monitor-stream-logs-to-event-hub

2. Azure function has the Event hub trigger and Cosmos output binding

a. Event Hub trigger for function

https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-event-hubs-trigger?tabs=csharp

QUESTION NO: 9

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

ANSWER: D

Explanation:

Azure Synapse Analytics is the best choice for hosting a data warehouse in Azure because it is designed for big data and data warehousing. It integrates with Azure Data Lake Storage Gen2, offers deep integration with Azure Machine Learning and other Azure services, and allows you to easily integrate and analyze all your data at scale using a single, unified environment. More information can be found on the official [Azure Synapse Analytics documentation](https://docs.microsoft.com/en-us/azure/synapse-analytics/).

QUESTION NO: 10

You have an on-premises application named App1 that uses an Oracle database.

You plan to use Azure Databricks to transform and load data from App1 to an Azure Synapse Analytics instance.

You need to ensure that the App1 data is available to Databricks.

Which two Azure services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Data Box Edge

B. Azure Data Lake Storage

C. Azure Data Factory

D. Azure Data Box Gateway

E. Azure Import/Export service

ANSWER: B C

Explanation:

To ensure the on-premises data from App1 is available to Azure Databricks for transformation and loading into Azure Synapse Analytics, you should use Azure Data Factory and Azure Data Lake Storage. Azure Data Factory is a data integration service that allows you to create data-driven workflows to orchestrate data movement and transformation. It can connect to on-premises Oracle databases and transfer the data to cloud services. Azure Data Lake Storage is a scalable and secure data lake for high-performance analytics workloads. Using these services will enable the seamless loading and

transformation of App1's data in Databricks before moving it to Synapse Analytics. For more details, visit the official [Azure Data Factory](https://docs.microsoft.com/en-us/azure/data-factory/introduction) and [Azure Data Lake Storage](https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-introduction) documentation.

QUESTION NO: 11 - (HOTSPOT)

You have an Azure Load Balancer named LB1 that balances requests to five Azure virtual machines.

You need to develop a monitoring solution for LB1. The solution must generate an alert when any of the following conditions are met:

Which signal should you include in the solution for each condition? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

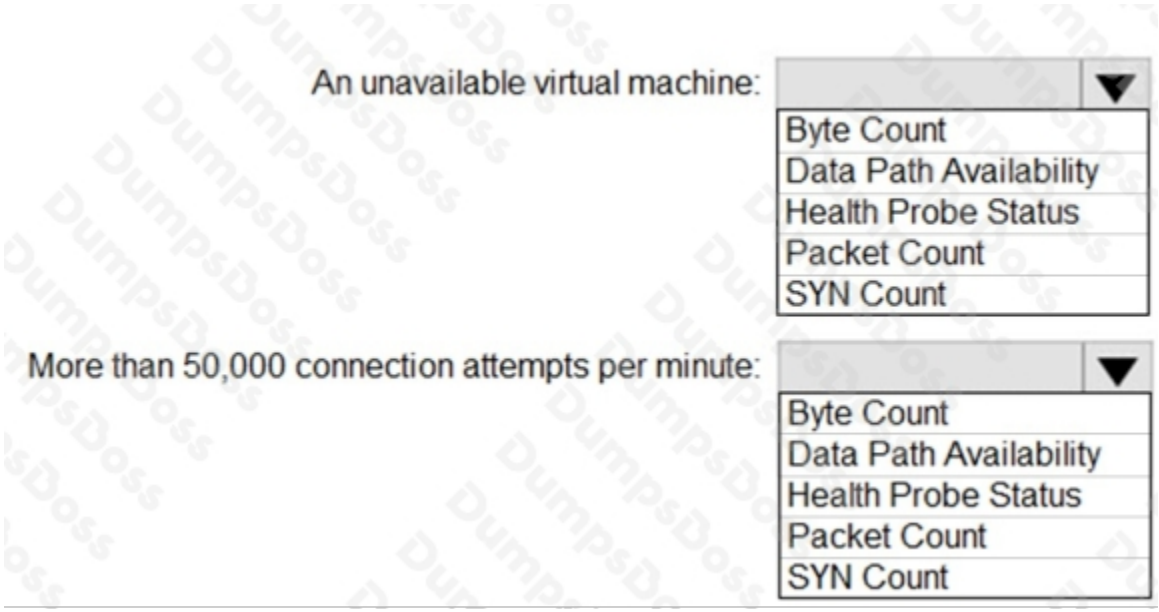

ANSWER:

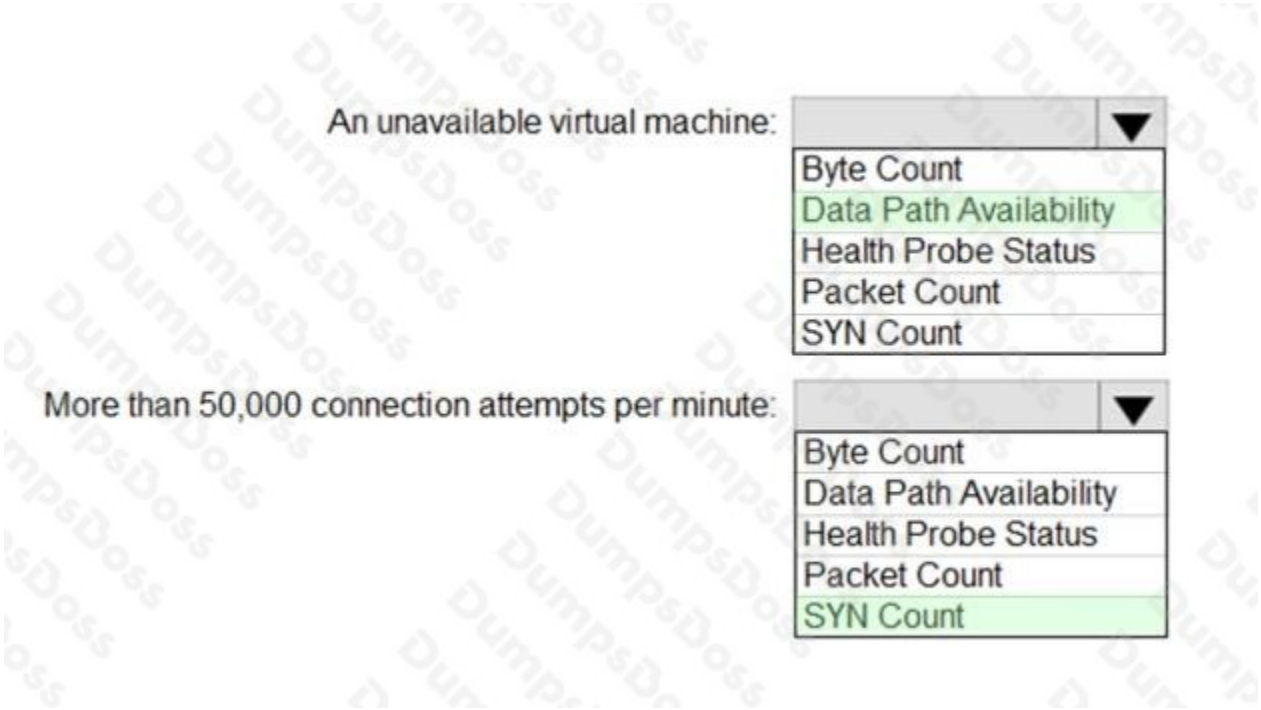

Explanation:

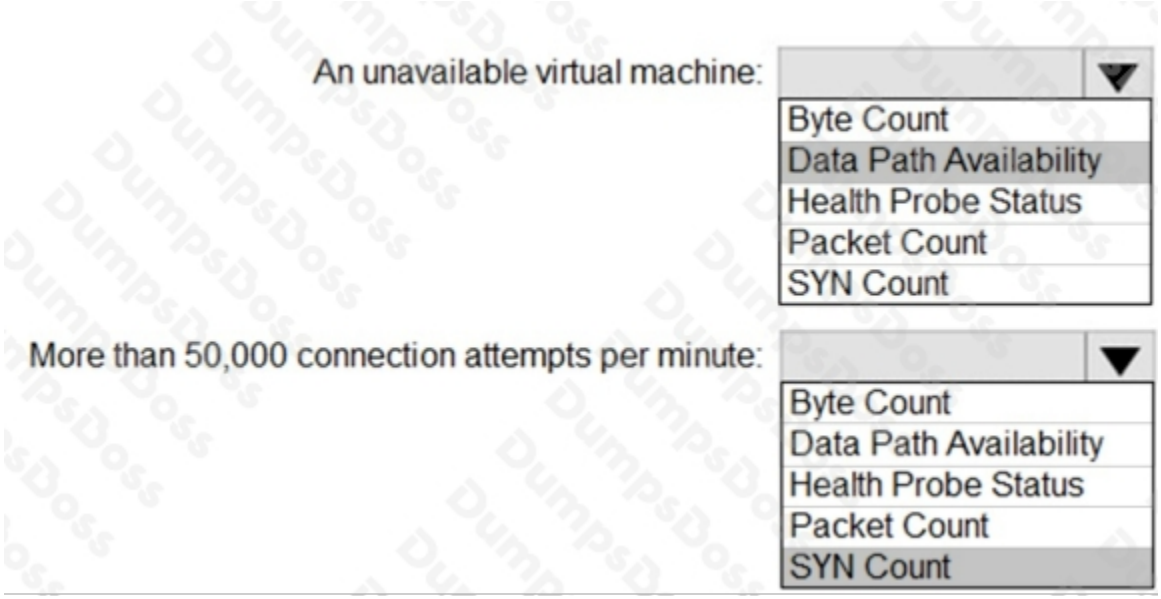

Box 1: Data path availability

Standard Load Balancer continuously exercises the data path from within a region to the load balancer front end, all the way to the SDN stack that supports your VM. As long as healthy instances remain, the measurement follows the same path as your application's load-balanced traffic. The data path that your customers use is also validated. The measurement is invisible to your application and does not interfere with other operations.

Note: Load balancer distributes inbound flows that arrive at the load balancer's front end to backend pool instances. These flows are according to configured load-balancing rules and health probes. The backend pool instances can be Azure Virtual Machines or instances in a virtual machine scale set.

Box 2: SYN count

SYN (synchronize) count: Standard Load Balancer does not terminate Transmission Control Protocol (TCP) connections or interact with TCP or UDP packet flows. Flows and their handshakes are always between the source and the VM instance. To better troubleshoot your TCP protocol scenarios, you can make use of SYN packets counters to understand how many TCP connection attempts are made. The metric reports the number of TCP SYN packets that were received.

Reference:

https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-standard-diagnostics

QUESTION NO: 12

The developers at your company are building a containerized Python Django app.

You need to recommend platform to host the app. The solution must meet the following requirements:

Which platform should you include in the recommendation?

A. Azure Container instances

B. an Azure App Service instance that uses containers

C. Azure Kubernetes Service (AKS)

ANSWER: C

Explanation:

To keep up with application demands in Azure Kubernetes Service (AKS), you may need to adjust the number of nodes that run your workloads. The cluster autoscaler component can watch for pods in your cluster that can't be scheduled because of resource constraints. When issues are detected, the number of nodes in a node pool is increased to meet the application demand.

Azure Container Registry is a private registry for hosting container images. It integrates well with orchestrators like Azure Container Service, including Docker Swarm, DC/OS, and the new Azure Kubernetes service.

Moreover, ACR provides capabilities such as Azure Active Directory-based authentication, webhook support, and delete operations.

Reference:

https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler

https://medium.com/velotio-perspectives/continuous-deployment-with-azure-kubernetes-service-azurecontainer-

registry-jenkins-ca337940151b

QUESTION NO: 13

You plan to store data in Azure Blob storage for many years. The stored data will be accessed rarely.

You need to ensure that the data in Blob storage is always available for immediate access. The solution must

minimize storage costs.

Which storage tier should you use?

A. Cool

B. Archive

C. Hot

ANSWER: A

Explanation:

Azure's Cool storage tier is specifically designed for infrequently accessed data which needs to be stored for at least 30 days, offering a cheaper alternative to the Hot tier while still enabling immediate access to data. The key advantage is cost savings on storage rather than on retrieval whenever such data is accessed correctly. It requires no special changes in your applications to access the data within this tier, making it ideal for backup, media files, and legal information which might not be needed frequently but should be immediately accessible when required.

For more information, refer to the official Azure Storage documentation.

QUESTION NO: 14

Your company has 300 virtual machines hosted in a VMware environment. The virtual machines vary in size and have various utilization levels.

You plan to move all the virtual machines to Azure.

You need to recommend how many and what size Azure virtual machines will be required to move the current workloads to Azure. The solution must minimize administrative effort.

What should you use to make the recommendation?

A. Azure Cost Management

B. Azure Pricing calculator

C. Azure Migrate

D. Azure Advisor

ANSWER: C

Explanation:

https://docs.microsoft.com/en-us/azure/migrate/migrate-appliance#collected-data---vmware

"Metadata discovered by the Azure Migrate appliance helps you to figure out whether servers are ready for migration to Azure, right-size servers, plans costs, and analyze application dependencies".

https://docs.microsoft.com/en-us/learn/modules/design-your-migration-to-azure/2-plan-your-azure-migration

QUESTION NO: 15 - (DRAG DROP)

Your company identifies the following business continuity and disaster recovery objectives for virtual machines that host sales, finance, and reporting application in the company's on-premises data center.

•The finance application requires that data be retained for seven years. In the event of a disaster, the application must be able to run from Azure. The recovery in objective (RTO) is 10 minutes,

• The reporting application must be able to recover point in-time data al a daily granularity. The RTO is eight hours.

•The sales application must be able to fail over to second on-premises data center.

You need to recommend which Azure services meet the business community and disaster recovery objectives. The solution must minimize costs.

What should you recommend for each application? To answer, drag the appropriate services to the correct application. Each service may be used owe. More than once not at an You may need to drag the spin bar between panes or scroll 10 view content.

ANSWER:

Explanation:

1) Sales: Azure Site Recovery only

2) Finance: Azure Site Recovery and Azure Backup

3) Reporting: Azure Backup only