Microsoft.AZ-104-KR.v2026-04-28.q180

□□□□:	AZ-104-KR
□□□□:	Microsoft Azure Administrator (AZ-104 Korean Version)
□□□:	Microsoft
□□ □□ □□□:	180
□□:	v2026-04-28
# □□ □:	772
# □□ □□□:	1800
https://www.krdump.com/Microsoft.AZ-104-KR.v2026-04-28.q180.html	

NEW QUESTION: 1
Azure Container Instances□ □□□□ Azure □□□ □□□□.
Azure □□□ □□□□□(CLI)□ Docker□ □□□ □□□□ □□□□.
image1□□□ □□□ □□□□ □□□□ □□□□.
□□□ Azure □□□□ □□□□□□ □□□□□□□ □□□□□□ image1□ □□□□ □□□.
□ □□ □□□ □□ □□ □□□ □□□□ □□□? □□□□ □□ □□□□ □□□ □□□□□. □□: □□ □□□ 1□□□□.
Answer Area

Answer:

Explanation:

This scenario requires two distinct tasks related to Azure Container Registry (ACR): creating the registry itself and uploading a container image into it. Microsoft Azure Administrator documentation clearly separates these responsibilities between Azure CLI and Docker CLI.

To provision a new Azure Container Registry, Microsoft documentation states that administrators must use the Azure CLI az acr create command. This command creates a private, managed Docker registry in Azure that can store and manage container images. The az acr create command defines the registry name, resource group, and pricing tier (Basic, Standard, or Premium). It is the only supported CLI command used to create an ACR resource. Commands such as az acr build, az container create, or docker create do not create a registry and therefore do not meet the requirement.

Once the registry exists, the container image must be uploaded to it. According to the Microsoft Azure Container Registry documentation, images are uploaded using standard Docker commands. After authenticating to the registry and tagging the image with the registry's login server name, the image is transferred using the docker push command. Microsoft explicitly states that Azure CLI does not upload local images directly to ACR; instead, Docker is responsible for pushing images.

The documented workflow is:

* Create the registry using az acr create.
* Authenticate Docker to the registry.
* Tag the image with the registry address.
* Upload the image using docker push.

Therefore, the only commands that correctly satisfy the requirements are:

* az acr create to provision the registry
* docker push to add image1 to the registry

Final Verified Answer:

* Provision a new container registry: az acr create
* Add image1 to the registry: docker push

NEW QUESTION: 2

User1□□□ □□□□ □□□ Azure □□□ □□□□.

User1□ □□ □□□ □□□□ □□ □□□□□ □□□ □ □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.

User1□ □□ □□ □□ □□□ □□(RBAC) □□□ □□□□ □□□?

A. □□□
B. □□ □□ □□□ □□□
C. □□□
D. □□ □□ □□□
Answer: D (LEAVE A REPLY)
To ensure that User1 can deploy virtual machines and manage virtual networks, you need to assign an RBAC role that grants the necessary permissions to perform these tasks. The solution must also use the principle of least privilege, which means that you should only grant the minimum level of access required to accomplish the goal.
Based on these requirements, the best RBAC role to assign to User1 is D. Virtual Machine Contributor. This role allows User1 to create and manage virtual machines, disks, snapshots, and network interfaces. It also allows User1 to connect virtual machines to existing virtual networks and subnets. However, it does not allow User1 to create or delete virtual networks or subnets, or to access the virtual machines themselves. This role follows the principle of least privilege by limiting User1's access to only the resources and actions that are relevant to deploying virtual machines and managing virtual networks1.

NEW QUESTION: 3
VM3□ □□□ □□ □□□ □□□□ □□□□ □□□ □□□□□□.
□□□ NSG□ □□□ □□□ □□□□ □□□.
□□□ □□□□ □□□?
A. VNet1□ □□□□□
B. Azure Advisor□ □□ □□ □□
C. Azure Monitor□ □□ □□
D. Traffic Manager □□□□ □□ □□ □ □□
E. Azure Network Watcher□□ IP □□ □□
Answer: E (LEAVE A REPLY)
Scenario: Litware must meet technical requirements including:
Ensure that VM3 can establish outbound connections over TCP port 8080 to the applications servers in the Montreal office.
IP flow verify checks if a packet is allowed or denied to or from a virtual machine. The information consists of direction, protocol, local IP, remote IP, local port, and remote port. If the packet is denied by a security group, the name of the rule that denied the packet is returned. While any source or destination IP can be chosen, IP flow verify helps administrators quickly diagnose connectivity issues from or to the internet and from or to the on-premises environment.
References:
https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-ip-flow-verify-overview

NEW QUESTION: 4
contoso.com□□□ Microsoft Entra □□□□ □□□□.
fabrikam.com□□□ □□ □□□□ □□□□□.
fabrikam.com□ □□□□ contoso.com □□□□ □□□□□ □□□.
□□□□ fabrikam.com □□□□□□ □□ □ □□□ □□ □□□.
Microsoft Entra □□ □□□□ □□□ □□ □□□?
A. □□ □□ □□□□ □□□ □□□ □□□ □□ □□□ □□□□□.
B. □□ □□□ □□□ □□□□ □□□ □□ □□□ □□□□□.
C. □□ □□ □□□□ □□ □□ □□□ □□□□□.
D. □□ □□□ □□□ □□□□ Microsoft □□□□ □□□ □□□□□.
Answer: C (LEAVE A REPLY)

Microsoft Entra ID provides External collaboration settings to control which external domains users can invite as guests. To restrict invitations so that only users from fabrikam.com can be invited, you must configure Collaboration restrictions.
Within the External collaboration settings, administrators can:
* Allow invitations only to users from specific domains
* Block invitations to all other external domains
Microsoft Entra documentation specifies that Collaboration restrictions are the control used to define allowed or blocked external domains for B2B guest invitations.
The other options do not meet the requirement:
* Guest user access restrictions control what guests can do after they are invited.
* Cross-tenant access - Tenant restrictions control inbound/outbound access behavior, not invitation eligibility.
* Microsoft cloud settings apply to cloud instances, not domain-based invitations.
Final Verified Answer:
C. From External collaboration settings, configure the Collaboration restrictions settings.

NEW QUESTION: 5
□□□□ □□□□ storage1□□□ Azure Storage □□□ □□□□.
□ □□□□ □□□ □□□ □□ □□□ □□□□ storage1□ □□ □□□□ □ □□□□ □□□□ □□□.
□ □□□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

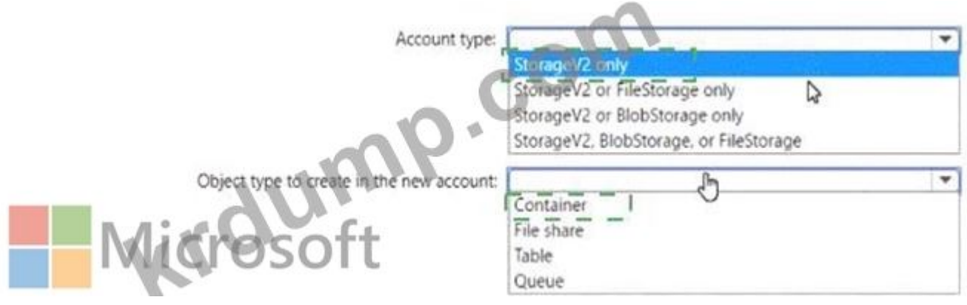

Explanation:
Account type: StorageV2 only
Object type to create in the new account: Container
To configure object replication in Azure Blob Storage, both the source and destination storage accounts must meet specific requirements. The Microsoft Azure Administrator documentation clearly specifies that object replication is supported only for StorageV2 (General-purpose v2) accounts, and replication can only occur between blob containers in those accounts.
* Account Type: Object replication is supported exclusively in StorageV2 accounts because older types such as BlobStorage or StorageV1 lack replication metadata and versioning support.

* Object Type: The replication applies only to blob containers-it does not support files, tables, or queues.
Replication works by asynchronously copying block blobs (not append or page blobs) between the source and destination containers once replication rules are configured.
Therefore, to replicate images (blobs) from storage1 to the new account, the new account must be created as StorageV2 only, and the object type must be a container.
Final Verified Answer:
Account type: StorageV2 only
Object type to create: Container

NEW QUESTION: 6
□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Location
Vnet1	US East
Vnet2	US East
Vnet3	US East
Vnet4	UK South
Vnet5	UK South
Vnet6	UK South
Vnet7	Asia East
Vnet8	Asia East
Vnet9	Asia East
Vnet10	Asia East

□□ □□ □□□□□ □□□□□ □□□□. □ □□ □□□□□□ 9□□ □□ □□□ □□□□.
Azure Boston□ □□□□ □□ □□□ □□ □□ RDP □□□ □□□□ □□□.
□□□ □ □□ □□ □□□ □□□□□?
A. 1
B. 3
C. 9
D. 10
Answer: B (LEAVE A REPLY)
According to the Microsoft documentation, Azure Bastion is a service that provides more secure and seamless RDP and SSH access to virtual machines without any exposure through public IP addresses. You can provision the service directly in your local or peered virtual network to get support for all the VMs within it.
In your scenario, you have three virtual networks that are peered with each other. This means that they can communicate with each other as if they were in the same virtual network. Therefore, you can deploy one Bastion host in any of the virtual networks and use it to connect to all the virtual machines in the peered virtual networks. You don't need to deploy a separate Bastion host for each virtual network or each virtual machine.
For more information about how to deploy and use Azure Bastion, see Tutorial: Deploy Bastion using specified settings: Azure portal.

NEW QUESTION: 7
Sub1□ Sub2□□ □ □□ Azure □□□ □□□, □□□ □□ □□□ Microsoft Entra □□□□ □□□□ □□□□.
□□ □□ □□ □□ □□ □□□□□ □□□□.

Name	Location	Subscription
VNet1	East US	Sub1
VNet2	East US	Sub1
VNet3	West US	Sub1
VNet4	East US	Sub2
VNet5	Central US	Sub2

Sub1□ Sub2□□ □ □□ Azure □□□ □□□, □□□ □□ □□□ Microsoft Entra □□□□ □□□□ □□□□.
□□ □□ □□ □□ □□ □□□□□ □□□□.
VNet1□ □□□□ □ □□ □□ □□□□□ □□□□□?
A. VNet2only
B. VNet2 □ VNet4□ □□
C. VNet2, VNet3, VNet4 □ VNet5
D. VNet2 □ VNet3□ □□
E. VNet2, VNet3 □ VNet4□ □□
Answer: (SHOW ANSWER)

NEW QUESTION: 8
Job1□□□ Azure Stream Analytics □□□ □□□ Azure □□□ □□□□.
Job1□ □□ □□ □□□□ □□□□□□ □□□□ □□ □□□ □□ □□□□ □□□.
□□ □□□ □□□□ □□□?
A. □□ □□□
B. □□□□ □□ □□□
C. □□□ □□□ □□□
D. □□ □□ □□□
Answer: B (LEAVE A REPLY)
Backlogged Input Events is a metric that shows the number of input events that are waiting to be processed by the Stream Analytics job1. This metric indicates the performance and health of the job, as well as the input data rate and latency. If the Backlogged Input Events metric is high or increasing, it means that the job is not able to keep up with the incoming events and some events are not processed in a timely manner2.
Output Events is a metric that shows the number of output events that are emitted by the Stream Analytics job1. This metric indicates the output data rate and throughput of the job. It does not show how many input events were not processed by the job.
Out-of-Order Events is a metric that shows the number of input events that arrive out of order based on their timestamp1. This metric indicates the quality and consistency of the input data source. It does not show how many input events were not processed by the job.
Late Input Events is a metric that shows the number of input events that arrive after the late arrival window has expired1. This metric indicates the timeliness and reliability of the input data source. It does not show how many input events were not processed by the job.

NEW QUESTION: 9
Group1□□□ □□□□ □□□ □□□ Azure □□□ □□□□. Group1□□ □□ □□ □□ □ □□ Azure □□□□ □□□□□ □□□□ □□□□.

Name	Resource request	Resource limit
container1	2 CPUs	2 CPUs
container2	3 CPUs	4 CPUs

container2□ container1□ □□□□ □□□ □□□ □□ CPU □□□□ □□□ □ □□□ □□□□ □□□.
□□□ □□ □□□?

A. □ □□□□□ □□□ □□□ □□□□□.
B. □□□□2□ □□□ □□□ 6□ CPU□ □□□□.
C. □□□□2□ □□□ □□□ 2□□ CPU□ □□□□.
D. □□□□1□ □□□ □□□ CPU 3□□ □□□□.
Answer: A (LEAVE A REPLY)

NEW QUESTION: 10
□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type
VM1	Virtual machine
storage1	Storage account
Workspace1	Log Analytics workspace
DB1	Azure SQL database

Azure Monitor□□ DCRI□□ □□□ □□ □□□ □□ □□□□□.
DCRI□□ □□ □□□□ □□□ □□□ □□□ □ □□, □□ □□□□ □□□□ □□□ □ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Data sources:
- VM1 only
- VM1 and storage1 only
- VM1, storage1, and DB1 only
- VM1, storage1, Workspace1, and DB1

Destinations:
- storage1 only
- Workspace1 only
- Workspace1 and storage1 only
- Workspace1, storage1, and DB1 only1

Answer:

Answer Area

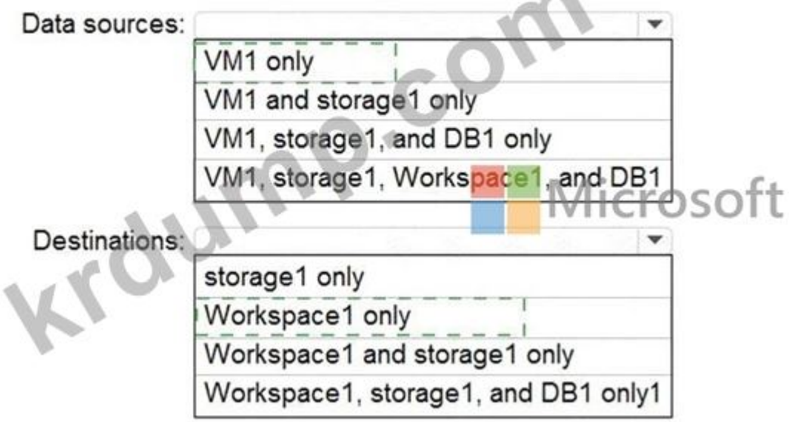

Explanation:
Data Sources: VM1 only
Destination: Workspace1 Only

NEW QUESTION: 11
□□ □□□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

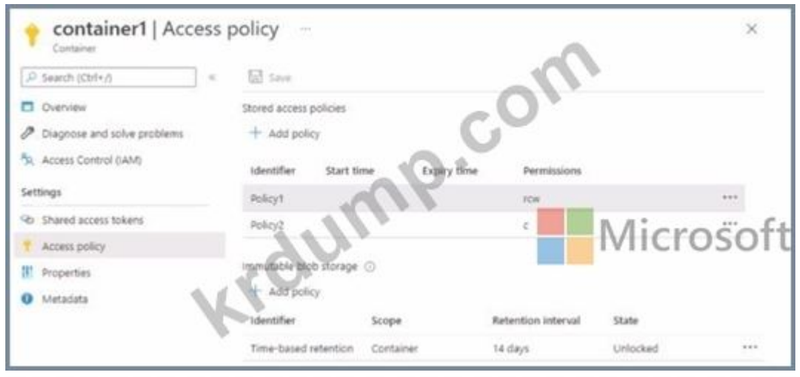

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

NEW QUESTION: 12

Azure □□□ □□□□.

□□ □□ □□□ □□□□ □□ □□ □□□ □□ □□□□□.

* □□□□ □□□□ □□□ □□ □□□□ □ □ □□□ □□□.

* □□□□ □□ □□□□□□ □□ □□□ □□□ □ □□□ □□□.

* □□□□ □□ □□□ □□□ □□□□ □□□.

□ □□ □□□ □□ □□ □□□□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

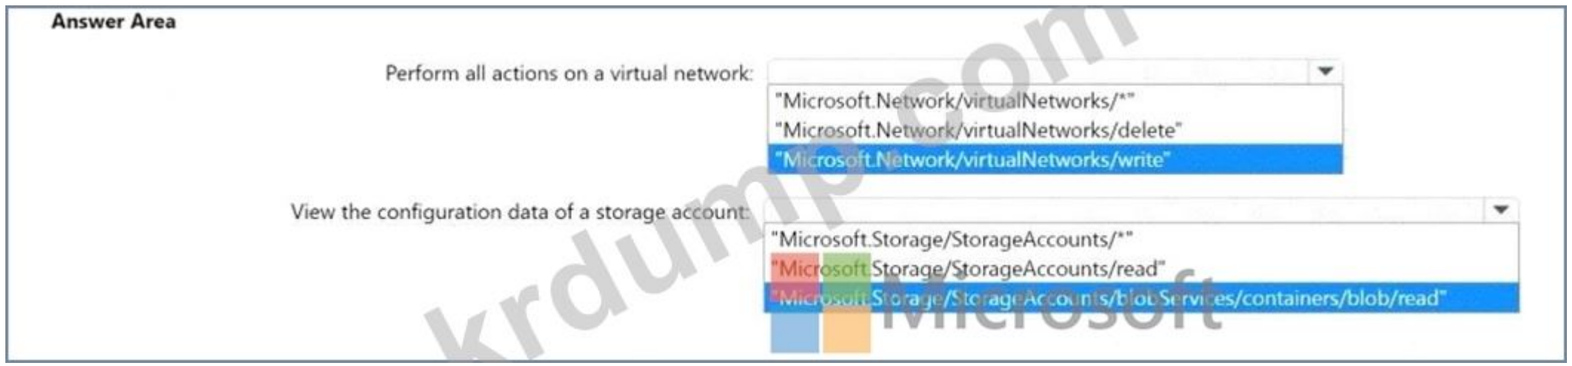

Answer:

Explanation:

Perform all actions on a virtual network:

"Microsoft.Network/virtualNetworks/*"

View the configuration data of a storage account:

"Microsoft.Storage/StorageAccounts/read"

To perform all actions on a virtual network, you need to use the wildcard () character in the action string, which grants access to all actions that match the string. The action string for virtual networks is "Microsoft. Network/virtualNetworks/". To view the configuration data of a storage account, you need to use the read action substring in the action string, which enables read actions (GET). The action string for storage accounts is "Microsoft.Storage/StorageAccounts/read". References:

https://learn.microsoft.com/en-us/azure/role-based-access-control/role-definitions

https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles

NEW QUESTION: 13

VM1□ VM2□□ □ □□ □□ □□□ □□□ Azure □□□ □□□□. Azure □□ □□ □□□ □□□□.

VM1□ VM2 □□ HTTPS □□□ □□□ □□□□ □□ □□ □□□ □□ □□□□□.

□□ □□ □□□ □□□□ □□ □□ □ □□ □□ □□ □□ □□□□ □□□□ □□□?

□ □□□ □□□□ □□□ □□□□□.

MOTL □□ 5□□ 1□□□□.

A. □□□□□ IP □□

B. □□□ □
C. □□ □□□
D. □□□□ NAT □□
E. □□ □□□□
Answer: A,C (LEAVE A REPLY)
To create a load balancing rule that will load balance HTTPS traffic between VM1 and VM2, you need to create two additional load balance resources: a frontend IP address and a health probe.
A frontend IP address is the IP address that the clients use to access the load balancer. It can be either public or private, depending on the type of load balancer. A frontend IP address is required for any load balancing rule1.
A health probe is used to monitor the health and availability of the backend instances. It can be either TCP, HTTP, or HTTPS, depending on the protocol of the load balancing rule. A health probe is required for any load balancing rule1.
A backend pool is a group of backend instances that receive the traffic from the load balancer. You already have a backend pool that contains VM1 and VM2, so you don't need to create another one.
An inbound NAT rule is used to forward traffic from a specific port on the frontend IP address to a specific port on a backend instance. It's not required for a load balancing rule, but it can be used to access individual instances for troubleshooting or maintenance purposes1.
A virtual network is a logical isolation of Azure resources within a region. It's not a load balance resource, but it's required for creating an internal load balancer or connecting virtual machines to a load balancer2.

NEW QUESTION: 14
□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type	Description
VNET1	Virtual network	Azure region: East US Contains the following subnets: • Subnet1: 172.16.1.0/24 • Subnet2: 172.16.2.0/24 • Subnet3: 172.16.3.0/24
VNET2	Virtual network	Azure region: West US Contains the following subnets: • DemoSubnet1: 172.16.1.0/24 • RecoverySubnetA: 172.16.5.0/24 • RecoverySubnetB: 172.16.3.0/24 • TestSubnet1:172.16.2.0/24
VM1	Virtual machine	Connected to Subnet2

Azure Site Recovery□ □□□□ □□ □□□ □□ □□ □□ □□ VM1□ □□□□□.
VM1□ □□□ □□□□□ □□□□ VNET2□ □□ □□ □□□□□ □□□□□.
VM1□ □□□ □□□ □□ □, □□ □□□ □□ □□□□ □□□□□?
A. □□□□□□1
B. □□□□□B
C. DemoSubnrt1
D. □□□□□A
Answer: A (LEAVE A REPLY)
https://learn.microsoft.com/en-us/azure/site-recovery/azure-to-azure-network-mapping The subnet of the target VM is selected based on the name of the subnet of the source VM.
- If a subnet with the same name as the source VM subnet is available in the target network, that subnet is set for the target VM.
- If a subnet with the same name doesn't exist in the target network, the first subnet in the alphabetical order is set as the target subnet.

NEW QUESTION: 15
□□ □□ □□ □□ □□□□ □□□ Sub1□□□ Azure □□□ □□□□.

Name	Description
RG1	Resource group
Action1	Action group that sends an email message to admin1@contoso.com

Sub1□□ □□□ □□ □□ □□□ □□□□ □□□□.
* □□: Alert1

* □□: Sub1□ □□ □□□ □□
o □□ □□ □□□□ □□□□□.
* □□ : □□ □□□□
* □□: Action1
Sub1□□ □□□ □□ □□ □□ □□□ □□□□ □□□□.
* □□: □□1
* □□: Sub1
* □□ □□: □□ □□
* □□ □□: □□ □□□
□□□ □□□: 2022□ 8□ 10□
o □□□: 2022□ 8□ 13□
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

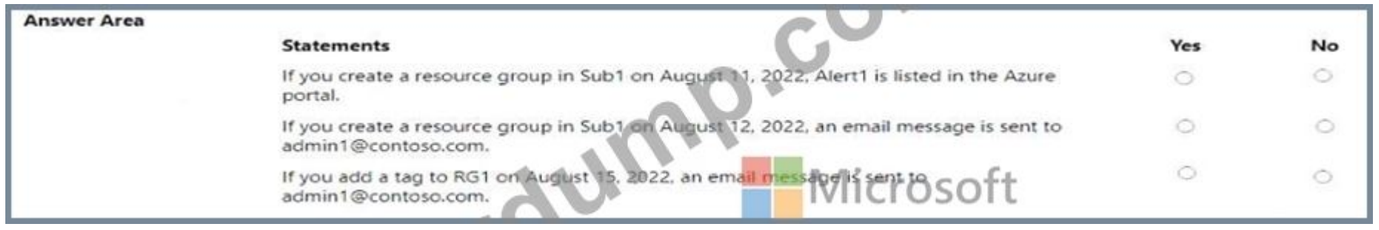

Answer:

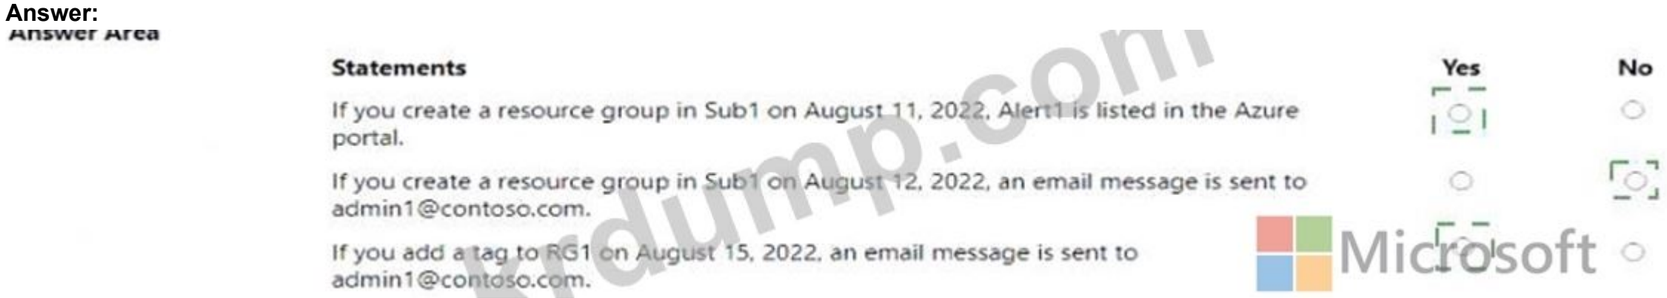

Explanation:

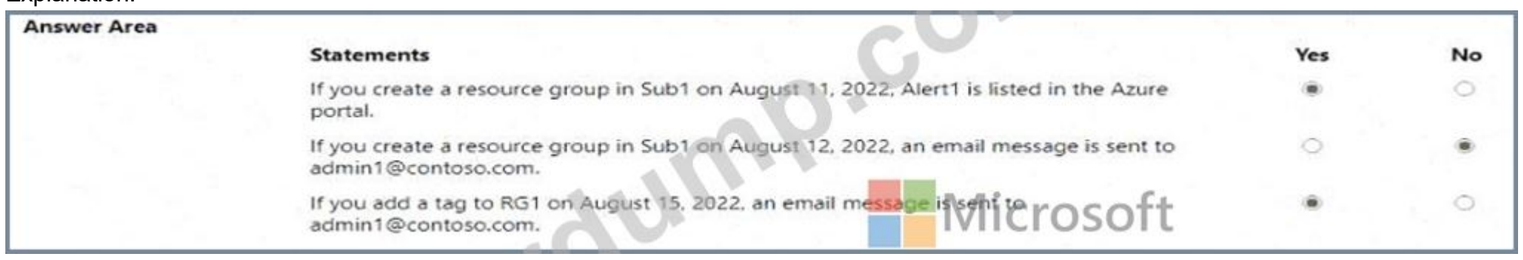

This question evaluates your understanding of Azure Monitor alerts, alert processing rules, and action groups.
1. Alert1 Configuration Details
Scope: All resource groups in Sub1
Include all future resources: # Yes
Condition: All administrative operations

This means any action that modifies Azure resources (create, delete, update, tagging, etc.) triggers the alert.
Action Group: Action1 # Sends email to admin1@contoso.com
So, Alert1 applies to all current and future resource groups in the subscription and triggers when any administrative operation occurs within that scope.
2. Alert Processing Rule (Rule1)
Scope: Sub1 (entire subscription)
Rule type: Suppress notifications
Start: August 10, 2022
End: August 13, 2022
According to Microsoft documentation:
"When a suppress notification rule is active, alert notifications (such as emails or SMS) are not sent, but the alerts themselves are still created and listed in Azure Monitor." So, between August 10-13, 2022, alerts will still appear in the Azure portal, but no email notifications are sent.
Now evaluate each statement:
Statement 1:
"If you create a resource group in Sub1 on August 11, 2022, Alert1 is listed in the Azure portal." The action (creating a resource group) triggers Alert1, since it's an administrative operation within scope.
However, August 11 is during the suppression period.
The alert is still created and visible in the Azure portal, but the notification is suppressed.
Answer: Yes (Alert appears in the portal even during suppression)
Statement 2:
"If you create a resource group in Sub1 on August 12, 2022, an email message is sent to admin1@contoso.
com."
Date: August 12, 2022 - still within the suppression period (Aug 10-13).
Therefore, no email notification is sent.
Answer: No
Statement 3:
"If you add a tag to RG1 on August 15, 2022, an email message is sent to admin1@contoso.com." Date: August 15, 2022, which is after the suppression period (ended Aug 13).
The alert condition ("All administrative operations") includes tag changes.
Therefore, this operation triggers an alert and sends an email through Action1.
Answer: Yes
Official Microsoft Azure Documentation Extract (AZ-104 Study Material Reference):
"Alert processing rules can temporarily suppress or redirect alert notifications without disabling or deleting the alert rule itself."
"When a suppress notification rule is active, alerts are still created in the Azure Monitor activity log, but notifications (emails, SMS, webhooks) are not sent."
"Administrative operations include all create, update, and delete actions at the resource and resource group level."

NEW QUESTION: 16
Plan1□□□ □□□ □□ Azure App Service □□□ □□□□.
□ □□ CPU □□□□ 80%□ □□□□ Plan1□ □□□□ □□□□□ □□ □□□. Plan1□ □□ □□□ □□□□ □□□?
A. □□□ □□ □□ □□□□ □□
B. □□□ □□ □□ □□□ □□ □□ □□
C. □□□□(□ □□□ □□) □□□ □□□□ P1
D. □□□□(□ □□□ □□) □□□ □□ S1
E. □□□ □□ □□ □□□ □□

Answer: B (LEAVE A REPLY)
Azure App Service Plans determine the scaling behavior of web apps hosted on them. Scaling can be done manually or automatically depending on the pricing tier.
From the Microsoft Azure Administrator Study Guide and official documentation ("Scale instance count manually or automatically" - Microsoft Learn):
"To enable automatic scaling based on metrics such as CPU usage, memory percentage, or HTTP queue length, your App Service Plan must be in the Standard, Premium, PremiumV2, or higher tier. You can then configure Scale Out (App Service plan) to use a Rules-Based method with performance thresholds." Available Scaling Options:
Manual: Fixed number of instances; no automatic scaling.
Automatic (Rules-Based): Create scaling rules based on metrics such as CPU > 80%, memory, or HTTP requests.
Scale Up (App Service Plan): Change pricing tier or hardware resources - does not provide auto-scaling.
In this scenario:
You already have a Standard plan (Plan1).
The requirement is to automatically scale out when CPU > 80%.
This behavior is achieved via Rules-Based scale-out, where you define a rule:
Metric: CPU Percentage
Condition: Greater than 80%
Action: Increase instance count
Therefore, you must choose Rules Based in the Scale out method settings for Plan1.

NEW QUESTION: 17
Azure □□□ □□□□.
□□ □□□ □□ □□□ □□ □□ □□ □□□ □□□□□.

□□□□ □□□ □□□□ □□□□ □□□ □□□ □□□□ □ □□□ □□ □□ □□□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer Area

At 9:00 AM, the scale set starts and CPU utilization is 80 percent for 15 minutes. How many virtual machine instances will be running at 9:15 AM?

- 2
- 3
- 4
- 5

At 10:00 AM, the scale set has five virtual machine instances running and CPU utilization falls to less than 15 percent for 60 minutes. How many virtual machine instances will be running at 11:00 AM?

- 1
- 2
- 3
- 4

Answer:

Answer Area

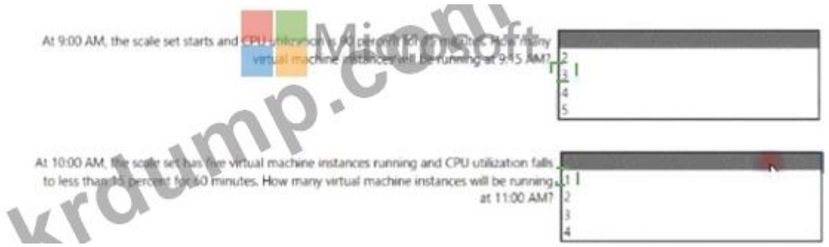

Explanation:

Box-1 : 3

Initial starts 2 VM's 15 minutes have passed. at 10 minutes 1 VM was added we now have 3 VM's. Cool down is 5 Minutes before another 10 minute wait cycle starts so the answer is 3.

Box-2: 1

Initial 5 VM's 60 minutes Pass. 1 VM removed every 15 minute cycle. 10 minutes wait timer plus 5 minute cool down equals 15 minutes cycle. Four 15 minute cycles pass equaling 60 minutes removing 4 VM's. We have 1 VM left.

Default Scale in and Out Default Durations are 10 minutes with 5 minute cool down.

The default scale set settings in Azure are:

-Minimum number of instances 1

-Maximum number of instances 10

-Scale out CPU threshold (%) 75

-Duration in minutes10

-Number of instances to increase by 1

-Scale in CPU threshold (%) 25

-Number of instances to decrease by -1

https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-autoscale- portal#create-a-rule-to-automatically-scale-in

NEW QUESTION: 18

VNet1□□□ □□ □□□□□ □□□ Azure □□□ □□□□. VNet1□ 10.0.0.0/16□ IP □□ □□□ □□□□ □□ □□ □□ □□ □□□□ □□□□□.

Name	IP address range
Subnet0	10.0.0.0/24
Subnet1	10.0.1.0/24
Subnet2	10.0.2.0/24
GatewaySubnet	10.0.254.0/24

Subnet1□□ □□□□ □□□□ VM1□□□ □□ □□□□□□□ □□□□ □□□□.

RT1□□□ □□□ □□□ □□□□ □□□□□.

□□ □□□□ □□□□ VM1□ □□ VNet1□ □□□□□ □□□.

RT1□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

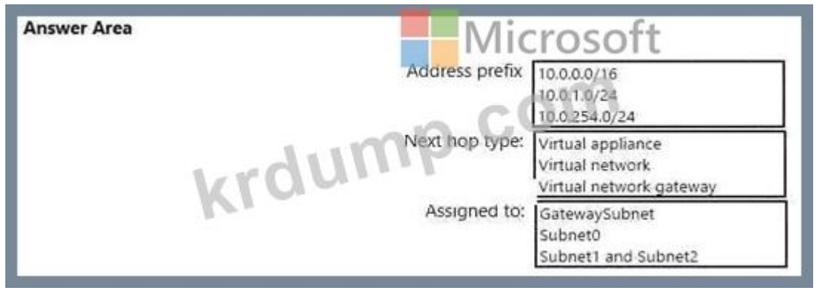

Answer:

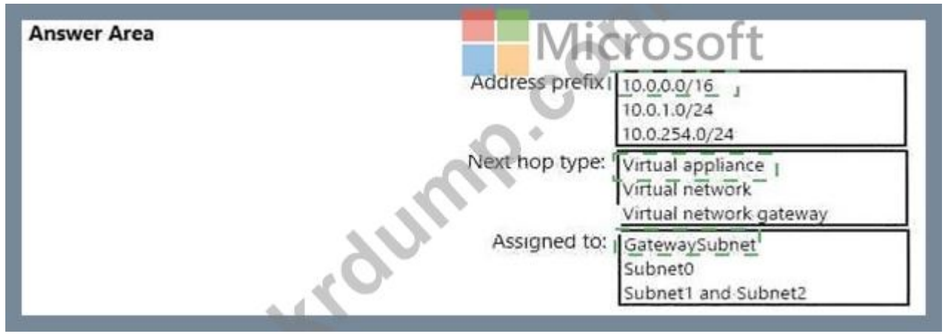

Explanation:

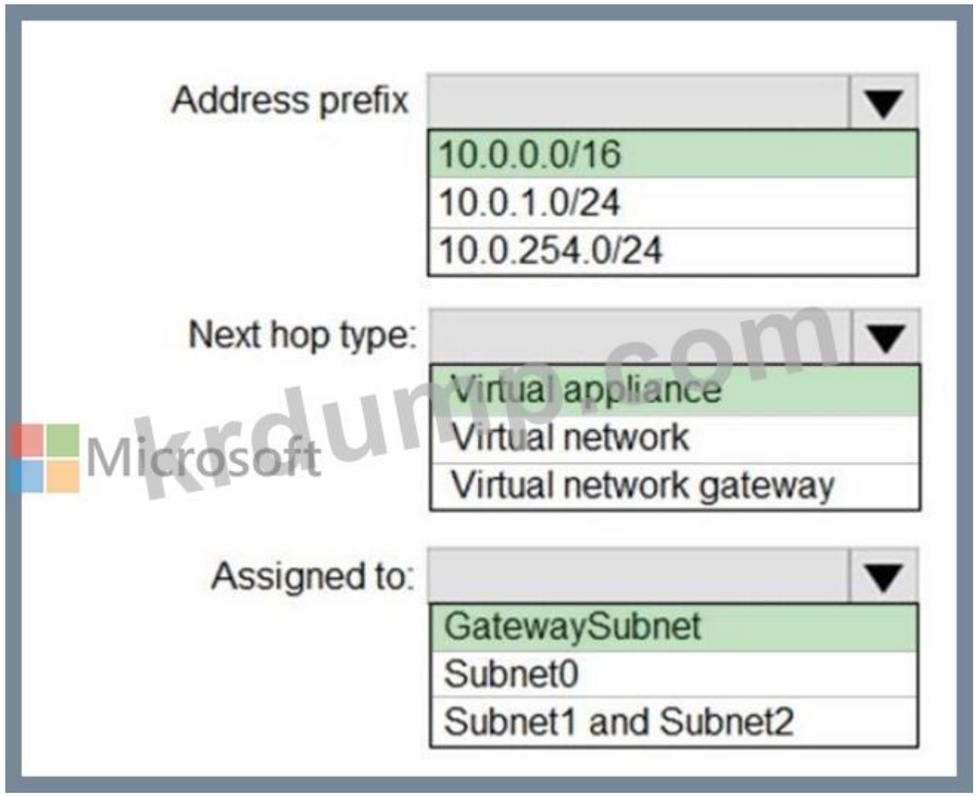

Box1 : 10.0.0.0/16

Address prefix in networking refer to the destination IP address range. In this scenario, destination is Vnet1 , hence Address prefix will be the address space of Vnet1.

Box 2 : Virtual appliance

Next hop gets the next hop type and IP address of a packet from a specific VM and NIC. Knowing the next hop helps you determine if traffic is being directed to the intended destination, or whether the traffic is being sent nowhere Next Hop --> VM1 --> Virtual Appliance (You can specify IP address of VM 1 when configuring next hop as virtual appliance) Box 3 : GatewaySubnet In the scenario it is asked for all the inbound traffic to Vnet1. Inbound traffic is flowing through SubnetGW.

You need to route all inbound traffic from the VPN gateway to VNet1 through VM1.So its traffic from Gateway subnet only.

Reference:

https://docs.microsoft.com/en-us/azure/virtual-network/manage-route-table#create-a-route-table

https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-next-hop-overview

NEW QUESTION: 19

□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

VNet1□ VNet2 □□ □□ □□□□ Microsoft □□ □□□□□ □□□□□ □□□□ □□□.

□□□ □□□□ □□□?

A. □□□□□□□

B. □□ □□□□□

C. □□□

D. □□ □□□

Answer: C (LEAVE A REPLY)
When you want to connect two Azure virtual networks (VNets) so that traffic between them stays on the Microsoft backbone network-and does not traverse the public internet-the correct configuration is Azure Virtual Network Peering.
Key Concept: Virtual Network Peering
According to Microsoft Azure Administrator Documentation (AZ-104 Study Guide & Azure Network Architecture Guide):
"Virtual network peering enables you to seamlessly connect two Azure virtual networks. The traffic between virtual machines in the peered networks uses the Microsoft backbone infrastructure and never traverses the public Internet, providing low latency, high bandwidth, and secure connectivity." Types of Peering:
VNet Peering (Intra-region) - connects VNets in the same Azure region.
Global VNet Peering - connects VNets across different Azure regions but within the same Azure cloud.
Both types use the Microsoft backbone network for communication, ensuring high-performance private connectivity.
Why Other Options Are Incorrect:
A). ExpressRoute:
ExpressRoute provides private connectivity between on-premises networks and Azure, not between Azure VNets. It is used for hybrid connectivity, not VNet-to-VNet communication inside Azure.
B). Private Endpoint:
Private Endpoints provide private access to Azure PaaS services (like Azure Storage, SQL Database, etc.), not for connecting two VNets directly.
D). Route Table:
Route tables (User-Defined Routes) control traffic flow within or between subnets/VNets, but they do not ensure traffic uses the Microsoft backbone. Without peering or a gateway, traffic would route via public IPs and the internet.
Official Microsoft Azure Documentation Extract:
From Microsoft Learn - Virtual Network Peering Overview:
"When virtual networks are peered, traffic between them is routed through the Microsoft backbone infrastructure, much like traffic between virtual machines in the same network. This ensures traffic never goes through the public Internet." Conclusion:
To guarantee that all traffic between VNet1 and VNet2 traverses Microsoft's private backbone network, you must configure VNet peering (either local or global depending on regions).

NEW QUESTION: 20
□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Location	Peered with
VNet1	East US	VNet2
VNet2	East US	VNet1

□ □□ □□□□□□ 50□□ □□□ □□ □□□ □□□□.
Azure Bastion□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
* □□□ □□□ □□□□□.
* □□ □□□ □ □□□□□ □□□□□.
* VNet1□ VNet2 □□□□ □□ □□□ □□□□□.
* Azure Bastion □□□□ □□ □□ □□□□□□.
Azure Bastion□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

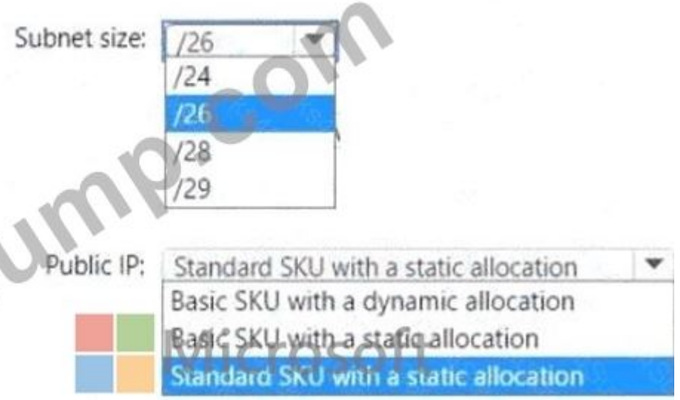

Answer:

Answer Area

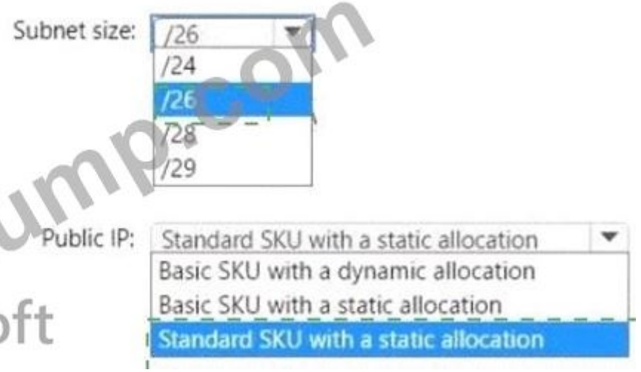

Explanation:

Azure Bastion is a fully managed service that provides secure and seamless RDP/SSH connectivity to your virtual machines directly through the Azure portal - without exposing those VMs to public IP addresses.

To meet the stated requirements, let's evaluate each configuration point using verified Azure documentation principles:

1## Support for host scaling:

Host scaling (auto-scale) is available only in the Standard SKU of Azure Bastion. The Basic SKU supports a single Bastion host instance and does not scale. Therefore, to support scaling, we must use the Standard SKU.

2## Support uploading and downloading files:

The file upload/download (RDP/SSH clipboard transfer) feature is supported only by the Standard SKU of Azure Bastion. The Basic SKU does not support these advanced capabilities.

3## Support for VMs in both VNet1 and VNet2:

Since VNet1 and VNet2 are in the same region (East US) and are peered, one Bastion host can be deployed in VNet1 and used to connect to VMs in both VNets. Cross-VNet connectivity for Bastion requires VNet peering and the Standard SKU.

4## Minimize the number of addresses on the Azure Bastion subnet:
Azure Bastion requires a dedicated subnet named AzureBastionSubnet.
* The minimum supported subnet size is /26 for the Standard SKU (as it supports scaling and multiple instances).
* The Basic SKU can use /27, but since we are using Standard SKU (for scaling and file transfer), the minimum possible subnet size is /26.This meets the requirement to minimize address space usage while supporting scaling.
5## Public IP requirements:
* The Standard SKU Bastion requires a Public IP address of SKU = Standard with Static allocation.
* Basic SKU Bastion can work with Basic Public IPs, but not Standard SKU Bastion.Hence, we must use a Standard SKU Public IP with Static allocation.
# Final Verified Configuration (per Microsoft Azure Administrator Documentation):
* Subnet size: /26
* Public IP: Standard SKU with a static allocation
Rationale Summary:
This configuration supports scaling, file transfer, cross-VNet connectivity, and minimal address consumption, satisfying all requirements as per official Azure documentation on Azure Bastion Standard SKU and Bastion network design guidelines.

NEW QUESTION: 21
□□□□□ □□□□□□ Share1□□□ SMB □□□ □□□□ □□□□.
□□ □□□□ □□□□ Azure □□□ □□□□.
webapp1□□□ □□
VNET1□□□ □□ □□□□
webapp1□ Share1□ □□□ □ □□□ □□□□ □□□.
□□□ □□□□ □□□?
A. Azure □□□□□□ □□□□□
B. Azure AD(Azure Active Directory) □□□□□□ □□□
C. Azure □□ □□□□ □□□□□
Answer: C (LEAVE A REPLY)
To enable a web app hosted in Azure App Service to connect securely to an on-premises SMB share (Share1), you must create hybrid network connectivity between your Azure environment and your on-premises network.
According to the Microsoft Azure Administrator Study Guide and Microsoft Learn documentation, Azure Web Apps running in an App Service Plan cannot directly access on-premises file shares over the public internet for security reasons. You must extend your on-premises network to Azure through a Virtual Network (VNet) and then integrate the web app with that network.
The Virtual Network Gateway is the component that enables this hybrid connectivity. It establishes a Site-to- Site VPN or ExpressRoute connection between the Azure VNet and the on-premises network, allowing the web app (after VNet integration) to access internal resources such as SMB shares, SQL Servers, or file servers.
Once the VPN gateway is configured and the web app is integrated with the VNet (Regional VNet Integration), the web app can securely access Share1 over the private network channel.
Official Microsoft Documentation Extract (Summary):
"To access on-premises resources from Azure App Service, configure VNet Integration and establish a Site-to- Site VPN or ExpressRoute connection using a Virtual Network Gateway. This allows Azure resources to securely communicate with on-premises systems such as SMB file shares or databases." (Source: Microsoft Learn - Connect an App Service app to an on-premises network using Azure VPN Gateway.)

NEW QUESTION: 22
Azure □□□ □□□□
□ □□□ □□□ □□□ □□□□□.
□□□ □□ □□□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
* □ □□□□ □□□ □□ □□ □ □□

* □□□□ □□ □□ □□□ □□□□□.
□□ □□□ □□ □□ □□ □□□ □□□□ □□□?

Answer:

Explanation:
RSA
4096

Key: RSA
length: 4096 https://learn.microsoft.com/en-us/azure/storage/common/customer-managed-keys-overview#key- vault-requirements

NEW QUESTION: 23
Windows Server 2019□ □□□□ VM1□□□ Azure VM□ □□□□.
VM1□ □□□ □□□ □□ □□□□□(□□ □□ □□□□□.)

Question | Exhibit

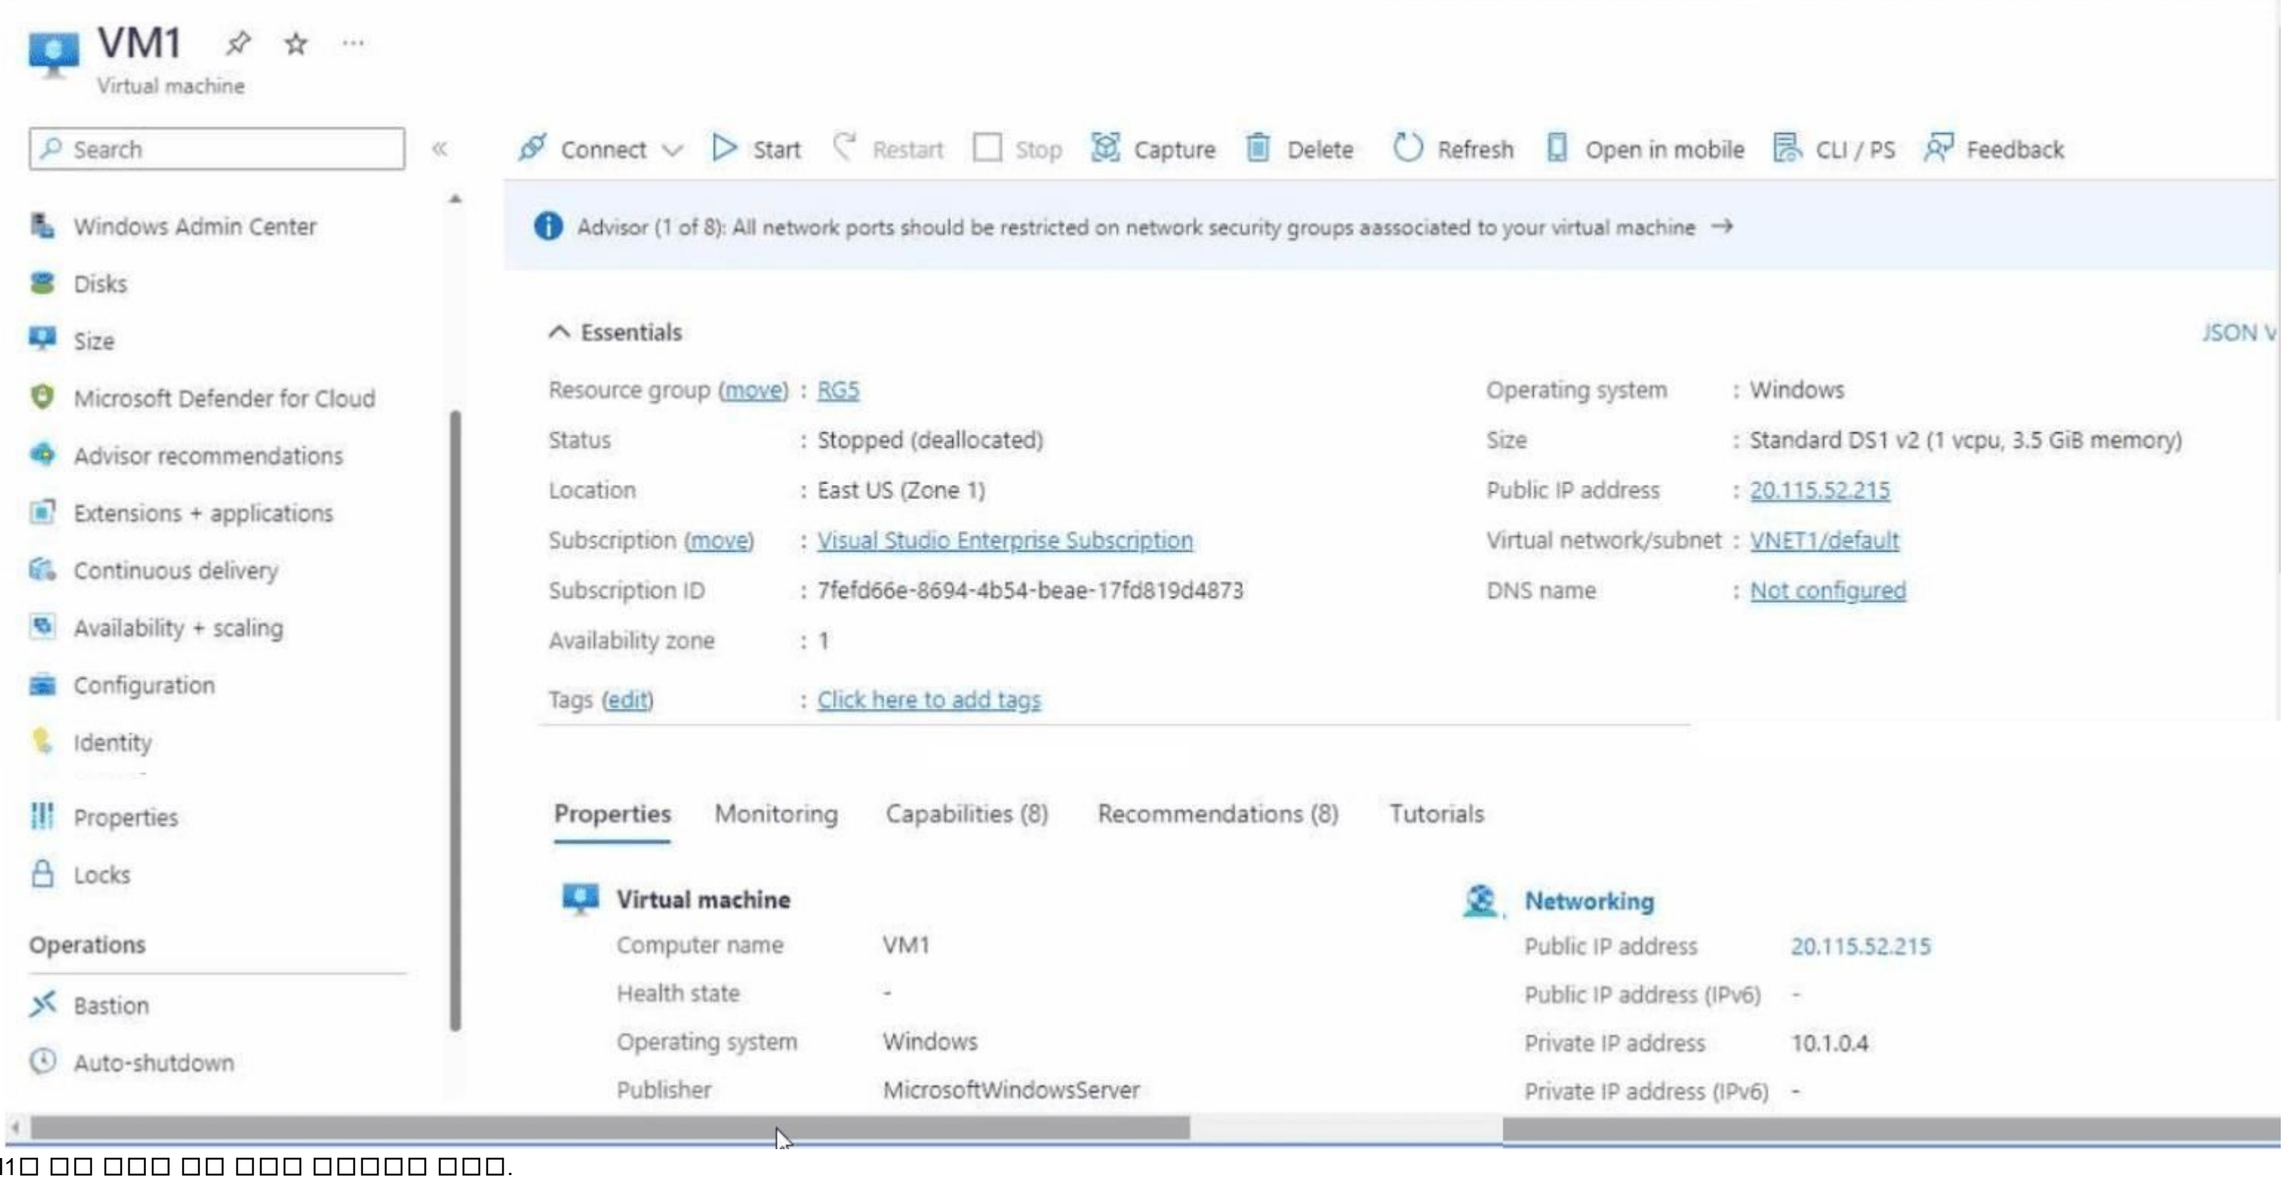

VM1□ □□ □□□ □□ □□□ □□□□□ □□□.

□□ □□ □□□ □□ □□□?

A. VM1□ □□ DNS □□□ □□□□□.

B. VM1□ □□□□□.

C. VM1□ □□□□ □□□□□.

D. VM1□ □□□□□.
Answer: (SHOW ANSWER)
Scenario Overview:
You have a Windows Server 2019 virtual machine (VM1) that has been created in Azure but is currently stopped (deallocated), as indicated by the exhibit (typical of this question from the AZ-104 pool). You need to enable Desired State Configuration (DSC) on this VM.
Understanding Azure Desired State Configuration (DSC)
Azure Desired State Configuration (DSC) is a management platform used to define and maintain configuration consistency for Azure virtual machines and other systems. In Azure, DSC can be managed via Azure Automation State Configuration, or locally using PowerShell DSC on the virtual machine.
To enable DSC on an Azure VM:
The VM must be running and accessible so that the DSC extension can be installed.
The Azure VM Agent must be operational (it runs only when the VM is started).
Configuration scripts or extensions (like Microsoft.Powershell.DSC) can only be applied to an active VM.
Why "Start VM1" is the Correct First Step
From the Microsoft Azure Administrator documentation (PowerShell DSC Extension for Azure VMs):
"Before you can enable the Azure Desired State Configuration extension on a virtual machine, the VM must be in the running state. The Azure VM Agent communicates with Azure to install and configure the DSC extension, and this agent only runs when the VM is powered on." This means that if VM1 is stopped (deallocated), Azure cannot deploy or enable the DSC extension, because it requires the VM to be active and connected to the Azure platform fabric.
Why the Other Options Are Incorrect
A). Configure a DNS name for VM1:
DNS naming allows external or internal connectivity but is not a prerequisite for enabling DSC. DSC configuration relies on the Azure VM Agent, not on DNS name resolution.
C). Capture a snapshot of VM1:
Snapshots are used for backup or replication, not for enabling configuration management features such as DSC.
D). Connect to VM1:
You cannot connect to VM1 (RDP/SSH) until it is started. Starting the VM is a prerequisite before any configuration management operation.
Key Documentation Reference (Microsoft Learn):
"The PowerShell Desired State Configuration (DSC) extension for Windows is used to configure Windows VMs. The extension requires the VM to be running so the Azure VM Agent can install and apply configurations."
- Source: Microsoft Learn - Azure Virtual Machine extensions and features overview / PowerShell DSC extension

NEW QUESTION: 24
Azure □□□ □□□□.
□□□ □□□□□ □□□ □□□ □□□□ □□□ □□□ □□□ □□□. Azure Monitor□□ □□ □□ □□□ □□□□ □□□ □□□□ □□□?
A. □□□, □□ □ □□ □□
B. □□□, □□ □ Microsoft 365 □□
C. Log Analytics □□ □□, □□□ □ □□ □□
D. □□□ □□ □□□□□, □□□□□□ □□ □□ □ □□□ □□
Answer: A (LEAVE A REPLY)
In Azure Monitor, Activity log alerts are used to detect when specific operations occur within your Azure subscription - such as creating, updating, or deleting resources, or in this case, removing a resource lock.
To configure an Activity Log alert, the following components are required:
* Scope (Resource or Subscription):Defines where the activity will be monitored - for example, an entire subscription, a specific resource group, or an individual resource.
* Condition:Specifies the event that triggers the alert. In this scenario, the event would be a "Delete management locks (Microsoft.Authorization/locks/delete)" operation.
* Action Group:Defines how the alert is delivered - for example, sending email notifications, SMS, push notifications, or invoking automation like Logic Apps or webhooks.

According to Microsoft Learn ("Create, view, and manage activity log alerts in Azure Monitor"):
"An activity log alert monitors a specific operation, at a specified scope, and sends notifications using an action group when that operation occurs." Microsoft 365 groups or data collection endpoints are not involved in Activity Log alerting. Similarly, Log Analytics workspaces are used for log queries, not activity log alerts.
Therefore, the correct configuration components to achieve the email alert are:
A resource, a condition, and an action group

NEW QUESTION: 25
contoso.onmicrosoft.com□□□ Azure Active Directory(Azure AD) □□□□ □□□□.
□□□ □□□ □□□ Admin1□□□ □□□□□ □□□□□.
□□ □□□□ user1@outlook.com □□□□ □□□□ Microsoft □□□ □□□ □□□□.
Admin1□ □□ □□□□ □□□□ Azure AD □□□□ □□□□□□ □□□□□ □□□ □□ □□ □□□□ □□□□□. "user1@outlook.com □□□□ □□□ □ □□□□. □□ □□ □□ □□□ □□□□□□." Admin1□ □ □ □□□□ □□□□ Azure AD □□□□ □□□□ □ □□□ □□□□ □□□.
□□□ □□ □□□?
A. □□ □ □□□ □□□□□□ Admin1□ □□ □□□ □□□ □□□□□.
B. □□ □□ □□□□□□ ID □□□□ □□□□□.
C. □□□ □□ □□□ □□ □□□□□□ □□□ □□ □□□□ □□□□□.
D. □□□ □□ □□□□□□ □□ □□ □□□ □□□□□.
Answer: D (LEAVE A REPLY)
In Azure AD, the ability to invite external (guest) users is controlled through External collaboration settings, found under Users # External collaboration settings in the Azure AD portal.
If a User administrator attempts to invite an external user but receives a "Generic authorization exception", this typically means the organization's external collaboration restrictions prevent user invitations.
According to Microsoft documentation (Configure external collaboration settings in Azure AD):
"To enable administrators or users to invite external guests, you must configure the External collaboration settings. You can define who can invite guest users into the directory, including administrators, users, or only specific roles." By updating these settings to allow User administrators to send invitations, Admin1 will be able to successfully invite the external partner (user1@outlook.com).

NEW QUESTION: 26
□□ □□ □□□ □□ □□□ □□□ Azure □□□ □□□□.

Name	Operating system	Connects to
VM1	Windows Server 2019	Subnet1
VM2	Windows Server 2019	Subnet2

VM1□ VM2□ □□ IP □□□ □□□□□. VM1□ VM2□ Windows Server 2019□□□ □□□□ □□ □□□□ □□□ □□□□□.
Subnet1□ Subnet2□ VNET1□□□ □□ □□□□□ □□□□.
□□□□ NSG1□ NSG2□□ □ □□ □□□□ □□ □□(NSG)□ □□□□ □□□□. NSG1□ □□ □□□ □□□□□.
NSG2□ □□ □□□ □□□ □□ □□□ □□ □□ □□□ □□□□□.
* □□□□: 100
* □□: □□1
* □□: 3389
* □□□□ : TCP
* □□ : Any

* □□□ : □□□
* □□: □□
NSG1□ □□□□ □□□□ □□, NSG2□ VM2□ □□□□ □□□□□□ □□□□ □□□□.
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
From the internet, you can connect to VM1 by using Remote Desktop.	○	○
From the internet, you can connect to VM2 by using Remote Desktop.	○	○
From VM1, you can connect to VM2 by using Remote Desktop.	○	○

Answer:

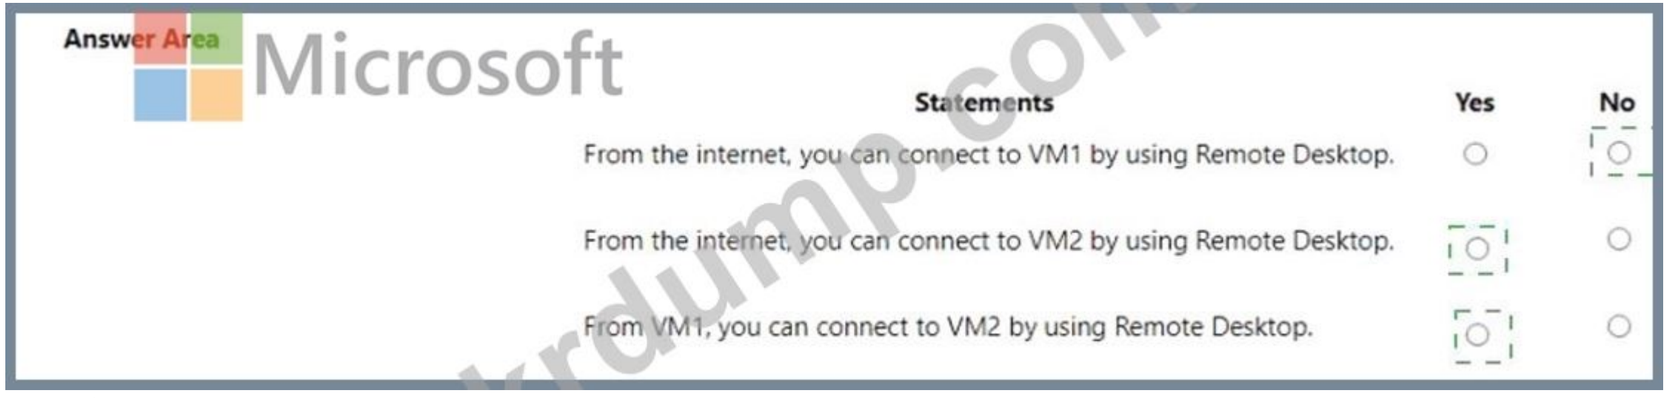

Explanation:

Answer Area

Statements	Yes	No
From the internet, you can connect to VM1 by using Remote Desktop.	○	●
From the internet, you can connect to VM2 by using Remote Desktop.	●	○
From VM1, you can connect to VM2 by using Remote Desktop.	●	○

In Azure, Network Security Groups (NSGs) control inbound and outbound traffic to network interfaces (NICs), subnets, and virtual machines (VMs) using rules based on priority and direction.
According to the Microsoft Azure Administrator Guide and Azure Networking documentation, the following principles apply:
Default NSG rules:
By default, an NSG denies all inbound traffic from the Internet except traffic originating from the same Virtual Network (VNet).
NSG allows all outbound traffic to the Internet.
Default inbound rules include:
Allow VNet inbound (priority 65000)
Allow Azure Load Balancer inbound (priority 65001)
Deny all inbound (priority 65500)
NSG associations:
NSGs can be associated either with a subnet or an individual network interface (NIC).
When both are applied, the NIC-level NSG takes precedence for inbound/outbound traffic.
If no explicit allow rule exists, default deny applies.
Analysis of Each Statement
1## From the internet, connect to VM1 by using RDP:
VM1's subnet (Subnet1) is associated with NSG1, which has only the default rules.
The default rules deny all inbound traffic from the Internet, including port 3389 (RDP).
Therefore, RDP from the Internet is blocked.
2## From the internet, connect to VM2 by using RDP:
VM2's NIC is associated with NSG2, which includes a custom allow rule (priority 100) permitting TCP traffic on port 3389 from any source to any destination.
This rule overrides the default deny rule.
Thus, RDP from the Internet is allowed.
3## From VM1, connect to VM2 by using RDP:
Both VMs reside in the same VNet (VNET1) but different subnets (Subnet1 and Subnet2).
The default NSG rule "Allow VNet Inbound" allows traffic between subnets within the same virtual network.
Therefore, VM1 can connect to VM2 via RDP (port 3389).

NEW QUESTION: 27
VM1□□□ □□ □□□ □□□ Azure □□□ □□□□.

VM1□ □□ □□ □□□ □□□□ 1TB □□□ □□□□ □□□ □□□□□.
* □□□ □□□ □□□□ □□□□ □□□ □□□□ □□□□□.
* □□ □□ □□ □□□ □□ □□ □□□ □□□□□.
* □□□□ □□□ □□□□ □□□ □□□ □□□□ □□□ □□□□□.
□ □□□ □□□□ □□ □□ □□□ □□□□□ □□□ □□□ □□□□ □□□□ □□□.

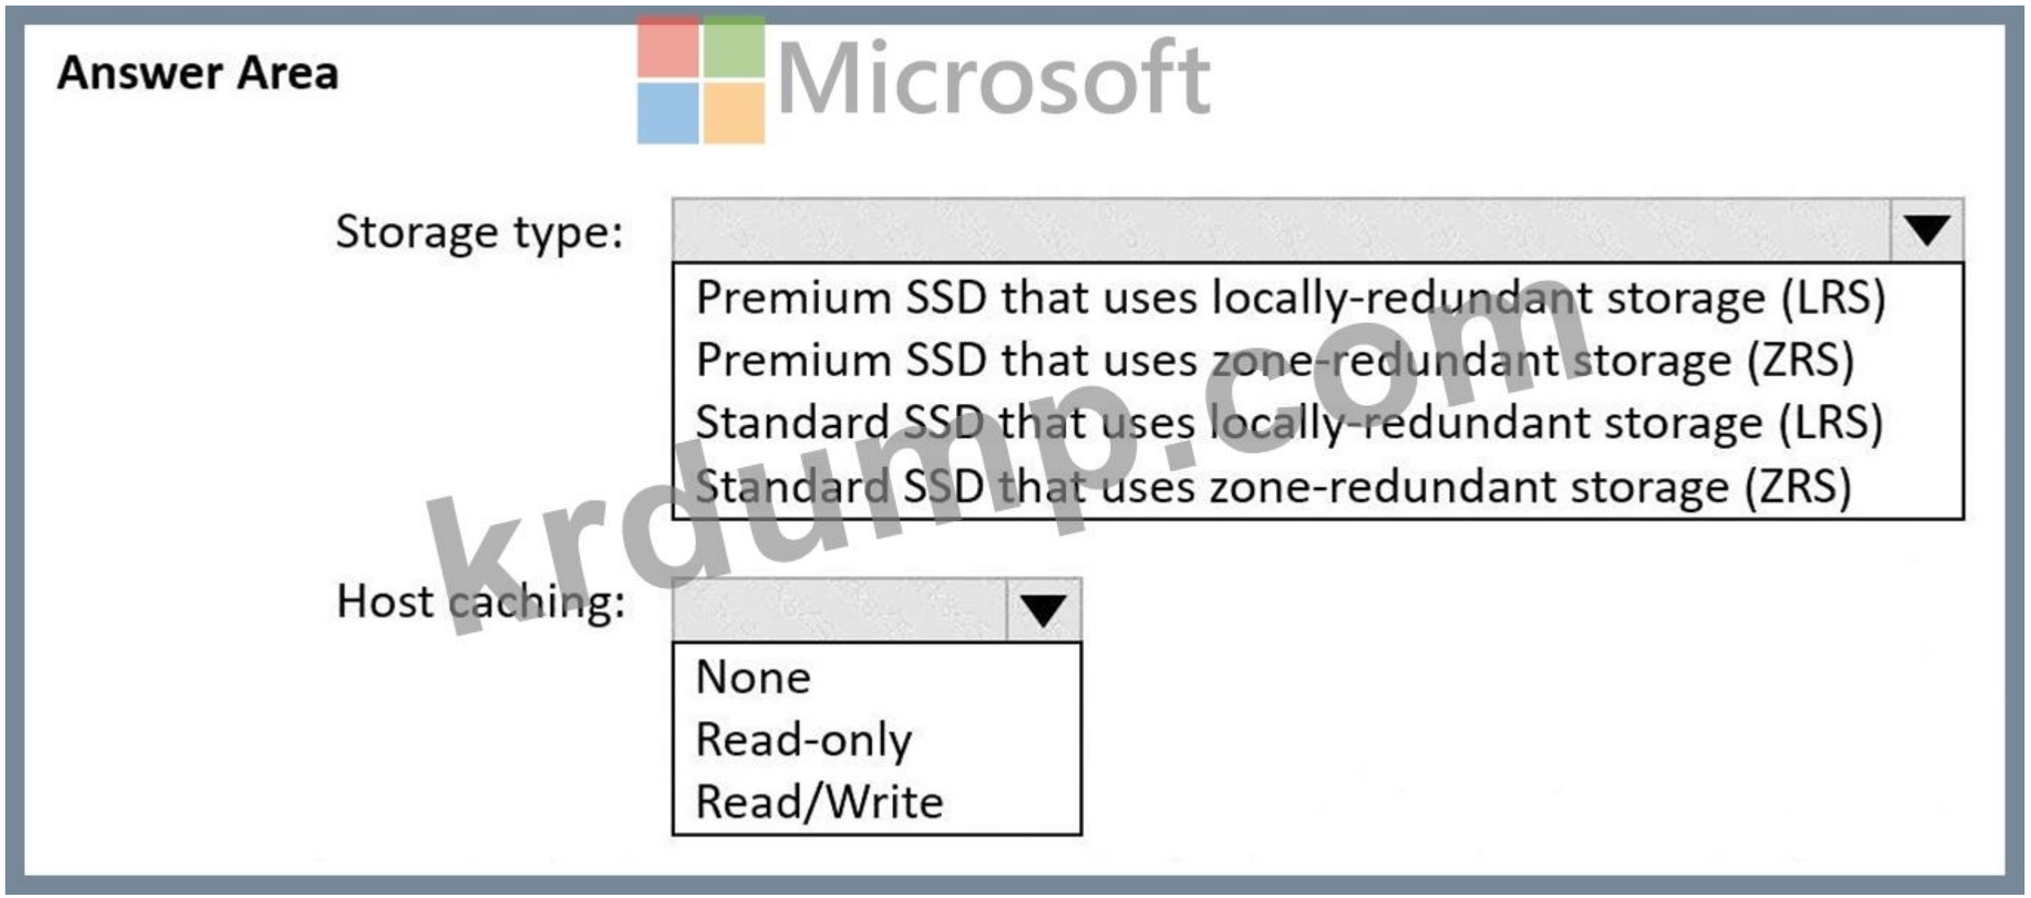

Answer:

Answer Area

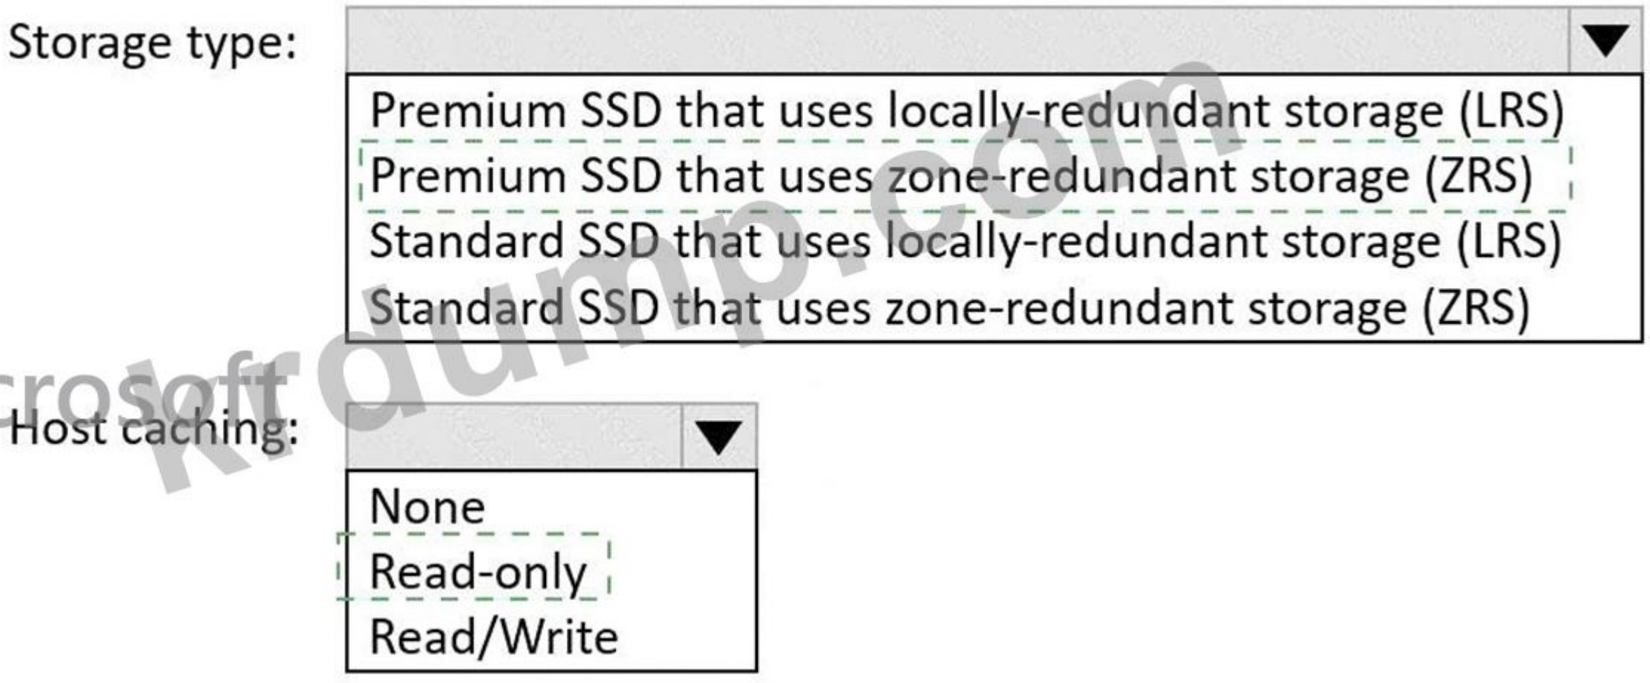

Explanation:
Storage Type: Premium SSD that uses zone-redundant storage (ZRS)
Host Caching: Read-only
The reasons for this recommendation are:
Premium SSD disks provide the lowest latency and the highest performance among the available disk types12.
Zone-redundant storage (ZRS) provides data resiliency in the event of a datacenter outage by replicating the data across three availability zones in the same region12.
Read-only host caching can improve the read performance of the disk by using the VM's RAM and local SSD as a cache13. This can also reduce the impact of a host failure on the disk data, as the cached data is not lost4.
Read/write host caching is not recommended for Premium SSD disks, as it can introduce additional latency and reduce the durability guarantees of the disk13.

NEW QUESTION: 28
NSG1□ NSG2□ □□ □□□ □□ □□□ □□□□□.
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
From VM1, you can establish a Remote Desktop session to VM2.	○	○
From VM2, you can ping VM3.	○	○
From VM2, you can establish a Remote Desktop session to VM3.	○	○

Answer:

Explanation:

NEW QUESTION: 29

RG1□□□ □□□ □□□ □□□□ Subscription1□□□ Azure □□□ □□□□.

RG1□□□ LB1□□□ □□ □□ □□□□ LB2□□ □□ □□ □□□□ □□□□.

Admin1□□□ □□□□ LB1□ LB2□ □□□ □ □□□ □□ □□□. □□□□ □□ □□ □□□ □□□ □□□.

□ □□□ □□ Admin1□□ □□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□□□: □□ 1□□ 1□□□□.

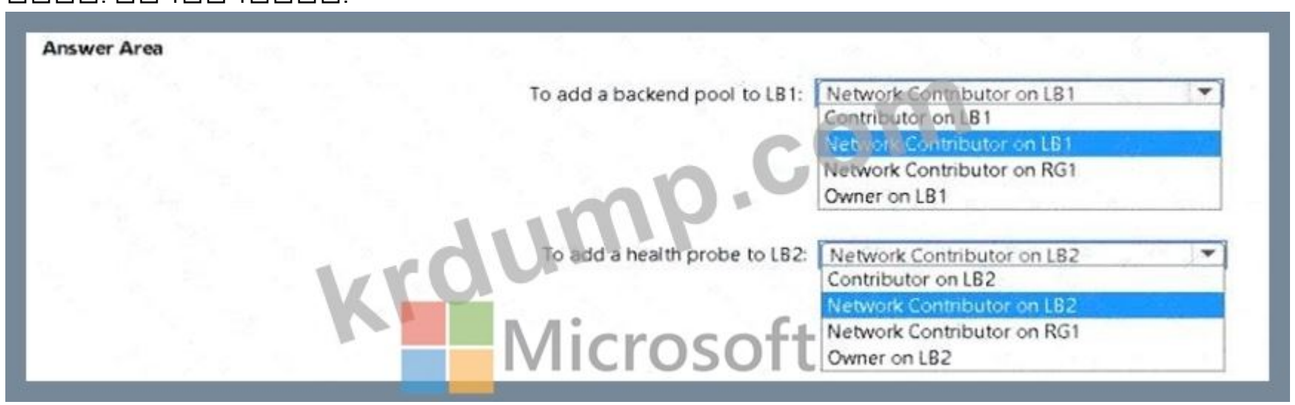

Answer:

Answer Area

Explanation:

This question tests your understanding of Azure RBAC (Role-Based Access Control) and the principle of least privilege when delegating permissions to manage specific load balancers.

Scenario Summary

You have:

Resource group: RG1

Two Load Balancers:

LB1 (Internal)

LB2 (Public)

You must allow Admin1 to manage configuration tasks on both load balancers individually:

Add a backend pool to LB1

Add a health probe to LB2

The goal is to assign the minimal required permissions (least privilege) needed for each operation.

Understanding the Role Requirements

1## Adding a Backend Pool to a Load Balancer

To add or modify a backend pool, you need permissions to:

Update the load balancer's configuration

Modify the associated NIC or VM backend association

The Network Contributor role includes these permissions.

Microsoft Learn - Network Contributor role permissions:

"Grants full access to manage network resources, including virtual networks, load balancers, network interfaces, and public IP addresses. Does not grant access to manage virtual machines or storage accounts." This means Network Contributor on LB1 (the load balancer resource itself) is sufficient to:

Add or remove backend pools

Configure load-balancing rules

Update frontend or backend associations

Correct Role: Network Contributor on LB1
No need for Contributor or Owner at the resource group level because that would grant more privileges than necessary (violates least privilege principle).
2## Adding a Health Probe to a Load Balancer
A health probe is a property of the load balancer resource itself.
To add or modify a health probe, you only need permissions to update the Load Balancer configuration.
Again, the Network Contributor role includes the required permissions to create or modify health probes.
Correct Role: Network Contributor on LB2
No additional access to RG1 or global resources is required.
Why Not Other Roles?
Role
Description
Too much / Too little
Contributor on RG1
Full access to all resources in the group
Too broad
Owner on LB1/LB2
Includes delete and permission management rights
Overprivileged
Network Contributor on RG1
Manage all network resources in RG1
Broader than needed
Network Contributor on LB1/LB2
Manage only the specified load balancer
Least privilege, correct
Final Verified Answers
Task
Role Assignment
To add a backend pool to LB1
Network Contributor on LB1
To add a health probe to LB2
Network Contributor on LB2
Microsoft Official Documentation Extracts (Azure RBAC & Load Balancer):
"The Network Contributor role can manage network resources but not grant access to others."
"To configure backend pools, health probes, and load balancing rules, assign the Network Contributor role on the load balancer resource itself."
"Follow the principle of least privilege by assigning resource-level roles instead of group- or subscription- level roles."
Final Verified Answer Summary:
Backend pool (LB1): Network Contributor on LB1
Health probe (LB2): Network Contributor on LB2

NEW QUESTION: 30
VM4□ □□ □□ □□ □□□ □□□□ □□□.

□□□ □□□ □□□□ □□□?

A. Azure □□ □□

B. Azure □□□ □□

C. Azure Logic □

D. Azure □□□ □□

Answer: B (LEAVE A REPLY)

Scenario: Create a workflow to send an email message when the settings of VM4 are modified.

You can start an automated logic app workflow when specific events happen in Azure resources or third-party resources. These resources can publish those events to an Azure event grid. In turn, the event grid pushes those events to subscribers that have queues, webhooks, or event hubs as endpoints. As a subscriber, your logic app can wait for those events from the event grid before running automated workflows to perform tasks - without you writing any code.

References:

https://docs.microsoft.com/en-us/azure/event-grid/monitor-virtual-machine-changes-event-grid-logic-app

NEW QUESTION: 31

□□ □□ □□ □□ □□□□ □□□ Sub1□□□ Azure □□□ □□□□.

Name	Description
RG1	Resource group
VNet1	Virtual network in RG1

□□ □□□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	Role	Role scope
Admin1	Owner	Sub1
Admin2	Contributor	RG1

Deploy.bicep□□□ □□□ Bicep □□□ □□□□.

```
param location string = resourceGroup().location

var virtualNetworkName = 'VNet2'
var subnetName = 'Subnet1'

resource virtualNetwork 'Microsoft.Network/virtualNetworks@2023-11-01' = {
  name: virtualNetworkName
  location: location
  properties: {
    addressSpace: {
      addressPrefixes: [
        '10.0.0.0/16'
      ]
    }
    subnets: [
      {
        name: subnetName
        properties: {
        addressPrefix: '10.0.0.0/24'
      }
    }
  ]
  }
}
```

□□ □□□ □□□□□.

```
New-AzResourceGroupDeploymentStack -Name Deploy1 -ResourceGroupName RG1 -TemplateFile Deploy.bicep -DenySettingsMode
DenyWriteAndDelete -ActionOnUnmanage DetachAll
```

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
Admin1 can delete VNet2.	○	○
Admin2 can add a subnet to VNet1.	○	○
Admin1 can add a subnet to VNet2.	○	○

Answer:

Answer Area

Explanation:
No
Yes
No
The command creates a resource group-scoped deployment stack in RG1 and enables deny settings using - DenySettingsMode DenyWriteAndDelete with -ActionOnUnmanage DetachAll. Deployment stacks manage the resources defined in the template (here, VNet2 and Subnet1) and can apply deny assignments that block changes to managed resources. Microsoft states that deny settings "define the operations that are prohibited on the managed resources" and that "this restriction applies to everyone unless they're explicitly granted access" (for example, excluded principals). Microsoft Learn Therefore, even though Admin1 is Owner at the subscription scope, the deny assignment still prevents delete and write operations against VNet2 outside controlled stack updates, so Admin1 cannot delete VNet2 and cannot add a subnet to VNet2 (adding a subnet is a write/update to the VNet). Microsoft Learn For VNet1, it already exists in RG1 and is not defined in the Bicep file, so it is not a managed resource of the stack. The deny settings apply to managed resources, so they do not block updates to VNet1. Admin2 has the Contributor role scoped to RG1, which Microsoft describes as granting "full access to manage all resources" at that scope (except assigning roles). Microsoft Learn Hence, Admin2 can add a subnet to VNet1.

NEW QUESTION: 32
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.
□□ □□ Azure □□□□ □□□□□ VNet1□□□ □□ □□□□□ □□□□□.
VNet1□ Windows Server□ □□□□ VM1 □ VM2□□ □ □□ □□ □□□ □□□□□□.
3□□ □□ VM1□□ VM2□□ □□ □□□□ □□□□ □□□□ □□□.
□□ □□: □□ □□□□□ DCS(□□□ □□□ □□)□ □□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
To inspect or capture network traffic between two Azure virtual machines (VM1 and VM2) within the same virtual network, Azure provides specialized network diagnostic tools - Azure Network Watcher, Connection Monitor, and Packet Capture - rather than Windows Performance Monitor.

The Performance Monitor (PerfMon) tool on Windows is used to collect metrics such as CPU, memory, and disk I/O performance counters, not network packet data or traffic analysis between VMs. While you could use PerfMon to track bandwidth usage or network throughput locally, it cannot capture or inspect packet-level communication or determine the content or flow between two Azure VMs.
According to the Microsoft Azure Administrator documentation (AZ-104 study guide), the correct method to analyze and inspect network traffic between Azure virtual machines includes:
Enabling Azure Network Watcher in the target region.
Using Packet Capture under Network Watcher to capture inbound and outbound packets.
Optionally using Connection Monitor to monitor connectivity metrics between endpoints.
Packet Capture allows you to define capture filters (e.g., source/destination IP, ports, protocols) and run it for a specific duration, such as three hours, as required in this scenario. The capture file is stored in Azure Storage or locally for later analysis with tools like Wireshark.
Therefore, creating a Data Collector Set (DCS) in Performance Monitor does not meet the requirement to inspect network traffic between VM1 and VM2.

NEW QUESTION: 33
□□ □□ □□ □□ □□ □□□□□ □□□□ Subscription1□□□ Azure □□□ □□□□.

Name	Subnet
VNet1	Sybnet11
VNet2	Subnet12
VNet3	Subnet13

Subscriptionon1□□ □□ □□ □□ □□□ □□□□ □□□□.

Name	IP address	Availability set
VM1	Subnet11	AS1
VM2	Subnet11	AS1
VM3	Subnet11	Not applicable
VM4	Subnet11	Not applicable
VM5	Subnet12	Not applicable
VM6	Subnet12	Not applicable

Subscription1□□ □□ □□□ □□ □□ □□□□ □□□□□.
□□: LB1
SKU: □□
□□: □□
□□□: Subnet12
□□ □□□□: VNET1
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Statements	Yes	No
LB1 can balance the traffic between VM1 and VM2.	○	○
LB1 can balance the traffic between VM3 and VM4.	○	○
LB1 can balance the traffic between VM5 and VM6.	○	○

Answer:

Statements	Yes	No
LB1 can balance the traffic between VM1 and VM2.	○	○
LB1 can balance the traffic between VM3 and VM4.	○	○
LB1 can balance the traffic between VM5 and VM6.	○	○

Explanation:

Yes

No

No

This question tests your understanding of Azure Resource Manager (ARM) template deployments and the copy loop function used to create multiple resources.

The provided ARM template is:

```
{
"$schema": "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
"contentVersion": "1.0.0.0",
"parameters": {},
"variables": {},
"resources": [
{
"type": "Microsoft.Resources/resourceGroups",
"apiVersion": "2018-05-01",
"location": "eastus",
"name": "[concat('RG', copyIndex())]",
"copy": {
"name": "copy",
"count": 4
}
}
],
"outputs": {}
}
```

Statement 1: The commands will create four new resources. # YES #

The ARM template includes a copy loop with "count": 4.

This means that the deployment will iterate four times and create four resource groups named sequentially as:

RG0

RG1

RG2

RG3
Each loop iteration generates one new resource group.
According to Microsoft Learn - Azure Resource Manager Template Documentation, the copy element is used to deploy multiple instances of a resource. Therefore, four resources will indeed be created.
Statement 2: The commands will create storage accounts in the West US Azure region. # NO # The ARM template resource type is clearly defined as:
"type": "Microsoft.Resources/resourceGroups"
This indicates that the deployment creates resource groups, not storage accounts.
Also, the "location" property in the template is explicitly set to "eastus".
Hence, the template creates four resource groups in the East US region, not in West US.
Therefore, this statement is false.
Statement 3: The first storage account that is created will have a prefix of 0. # YES # The name of each resource is defined by:
"name": "[concat('RG', copyIndex())]"
The function copyIndex() starts with a zero-based index, meaning the first iteration uses 0.
So, the first resource created will be:
RG0
This behavior is documented in Azure ARM Templates - copyIndex() function, which specifies that the default starting index is 0 unless an offset value is provided.
Therefore, this statement is true.
Final Verified Answers:
Statement
Answer
Explanation Summary
The commands will create four new resources.
Yes
The copy loop count = 4
The commands will create storage accounts in the West US region.
No
Creates resource groups in East US
The first storage account that is created will have a prefix of 0.
Yes
copyIndex() starts at 0 # RG0

NEW QUESTION: 34
VNET2□ □□ □□□□ □□ □□□ □□ □□□□□.

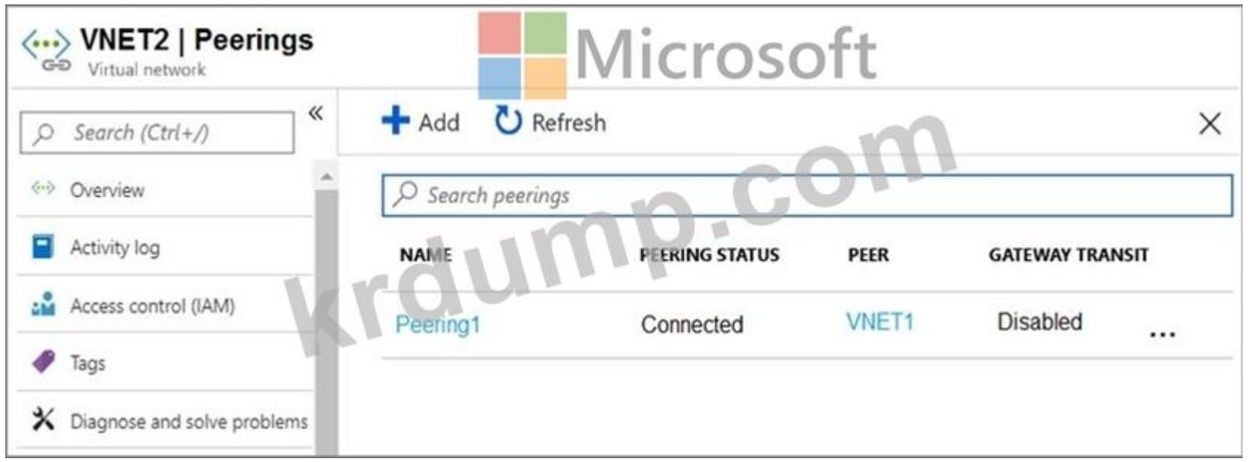

VNET3□ □□ □□□□ □□ □□□ □□ □□□□□.

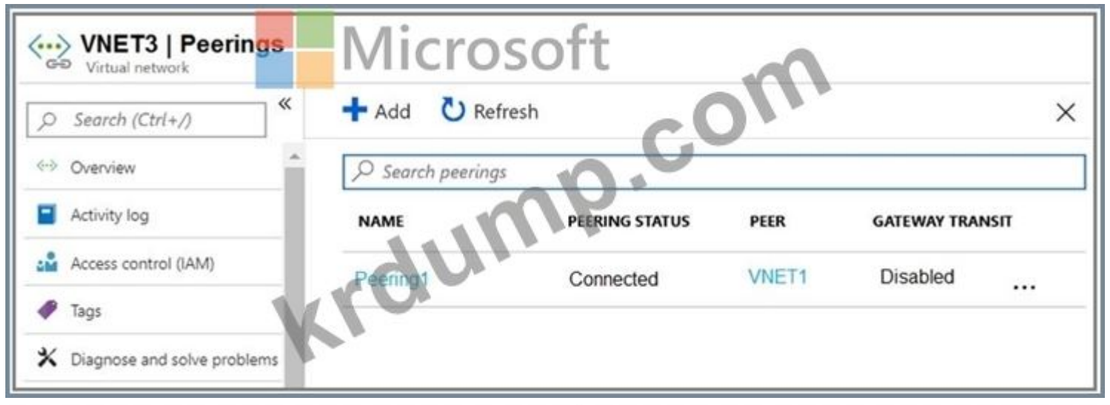

□□ □□□□ □□ □□□ □□□ □□□□ □ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

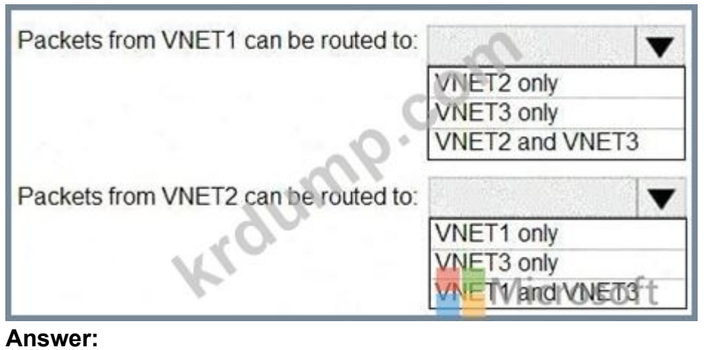

Answer:

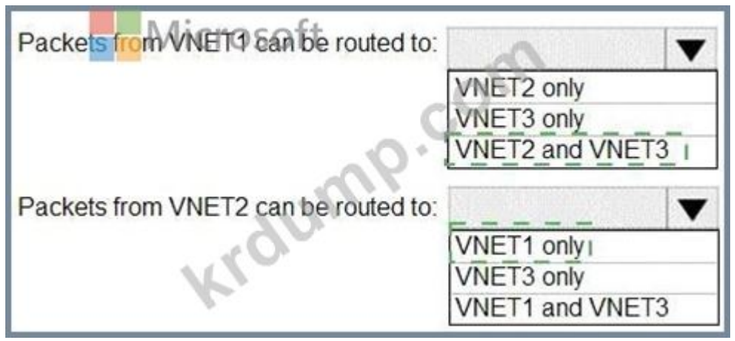

Explanation:

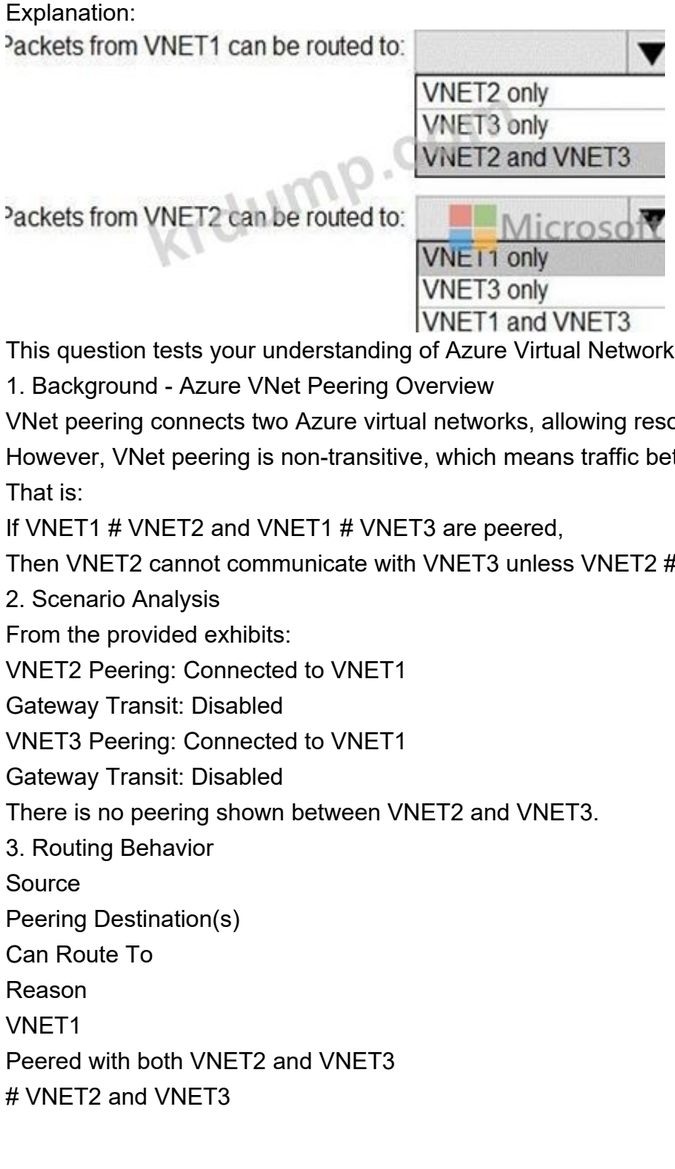

This question tests your understanding of Azure Virtual Network (VNet) Peering and its transitivity limitations.

1. Background - Azure VNet Peering Overview

VNet peering connects two Azure virtual networks, allowing resources in those VNets to communicate with each other over Microsoft's private backbone network.

However, VNet peering is non-transitive, which means traffic between two VNets can only flow directly between them if each VNet is peered explicitly.

That is:

If VNET1 # VNET2 and VNET1 # VNET3 are peered,

Then VNET2 cannot communicate with VNET3 unless VNET2 # VNET3 peering also exists.

2. Scenario Analysis

From the provided exhibits:

VNET2 Peering: Connected to VNET1

Gateway Transit: Disabled

VNET3 Peering: Connected to VNET1

Gateway Transit: Disabled

There is no peering shown between VNET2 and VNET3.

3. Routing Behavior

Source

Peering Destination(s)

Can Route To

Reason

VNET1

Peered with both VNET2 and VNET3

VNET2 and VNET3

VNET1 has direct peer connections to both VNets. Packets from VNET1 can reach both.
VNET2
Peered only with VNET1
VNET3 (no direct peering)
Azure VNet peering is non-transitive - packets cannot be forwarded via VNET1 to VNET3. Therefore, VNET2 can only send traffic to VNET1.
4. Microsoft Documentation Extract (Azure Official Docs)
"VNet peering is non-transitive. If VNetA is peered with VNetB, and VNetB is peered with VNetC, VNetA cannot automatically communicate with VNetC unless a direct peering between VNetA and VNetC is also established." (Source: Microsoft Learn - "Create, change, or delete a virtual network peering" and "Virtual network peering overview") Also:
"Gateway transit allows one peered network to use another's VPN gateway, but does not affect the basic non- transitive nature of peering." Since Gateway Transit = Disabled, this feature does not apply here.
5. Final Routing Summary
From
To
Routing Allowed
Explanation:
VNET1 # VNET2
Yes
Direct peering exists
VNET1 # VNET3
Yes
Direct peering exists
VNET2 # VNET3
No
No direct peering; peering is not transitive
VNET2 # VNET1
Yes
Direct peering exists
Final Verified Answer:
Packets from VNET1 can be routed to: VNET2 and VNET3
Packets from VNET2 can be routed to: VNET1 only
Microsoft Azure Administrator Study Guide - Official Reference Summary:
"VNet peering enables full mesh connectivity but does not automatically enable transitive routing."
"To establish cross-VNet communication, you must create a direct peering between each pair of VNets."
"Disabling Gateway Transit prevents shared routing or gateway propagation." (Reference: Microsoft Learn # Azure Virtual Network Peering Overview, AZ-104 Exam Objective: Configure and manage virtual networking)

NEW QUESTION: 35
Azure □□□ □□□□.
□□ □□□ □□ Azure Storage □□□ □□ □□□□□.


Microsoft Azure
Search resources, services, and docs (G+/)
Home > Subscriptions > Subscription1 - Resources > New > Create storage account
Create storage account
Validation passed
Microsoft
Basics Networking Advanced Tags Review + create
Basics
Subscription Subscription1
Resource group RG1
Location (Europe) North Europe
Storage account name storage16852
Deployment model Resource manager
Account kind StorageV2 (general purpose v2)
Replication Locally-redundant storage (LRS)
Performance Standard
Access tier (default) Hot
Networking
Connectivity method Private endpoint
Private Endpoint (New) StorageEndpoint1 (blob) (privatelink.blob.core.windows.net)
Advanced
Secure transfer required Enabled
Large file shares Disabled
Blob soft delete Disabled
Blob change feed Disabled
Hierarchical namespace Disabled
NFS v3 Disabled
Create
< Previous
Next >
Download a template for automation

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

Azure Storage ensures durability and high availability of data by replicating it within and/or across datacenters depending on the replication option selected. The exhibit shows that the replication type configured is Locally-redundant storage (LRS), and the access tier (default) is Hot.

1. Minimum Number of Copies - Locally-redundant storage (LRS)

According to Microsoft Azure Storage Documentation (AZ-104 Study Guide):

"Locally redundant storage (LRS) replicates your data three times (3 copies) within a single physical location in the primary region." Each piece of data is written synchronously to three separate storage nodes in the same datacenter. This provides protection against hardware failures within that facility.

Replication options and their redundancy levels:
Replication Type
Number of Copies
Location of Copies
Locally-redundant storage (LRS)
3
Same datacenter (single region)
Zone-redundant storage (ZRS)
3
Across availability zones in same region
Geo-redundant storage (GRS)
6
3 copies in primary region + 3 in paired region
Read-access geo-redundant storage (RA-GRS)
6
Same as GRS, plus read access to secondary region
Therefore, with LRS, the minimum number of copies is 3.
2. Reducing Cost of Infrequently Accessed Data
Azure Storage provides access tiers designed for different usage patterns:
Tier
Description
Typical Use Case
Hot
Highest storage cost, lowest access cost
Frequently accessed data
Cool
Lower storage cost, higher access cost
Infrequently accessed data
Archive
Lowest storage cost, highest retrieval cost
Long-term, rarely accessed data
As per Microsoft Learn - Manage access tiers in Azure Blob Storage:
"To reduce costs for data that is infrequently accessed, you can move blobs from the Hot tier to the Cool or Archive access tier." Hence, to minimize the cost of infrequently accessed data, you should modify the Access tier (default) setting from Hot to Cool or Archive.
Official Microsoft Extract:
From Microsoft Learn - Redundancy in Azure Storage:
"With LRS, three copies of your data exist in a single datacenter."
"To optimize storage costs for infrequently accessed data, set the access tier to Cool or Archive."
Final Verified Answers:
Statement
Correct Answer

The minimum number of copies of the storage account will be:

3

To reduce the cost of infrequently accessed data in the storage account, you must modify the:

NEW QUESTION: 36

□□ Azure Resource Manager □□□□ □□□□ Azure □□□□ □□□□□ □□□ □□□□□.

□□□□ □□□ □□□□ □□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□.

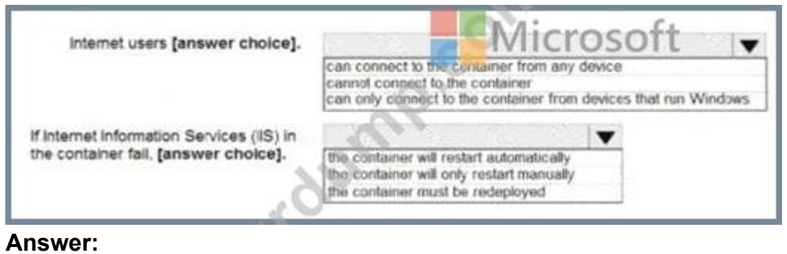

Answer:

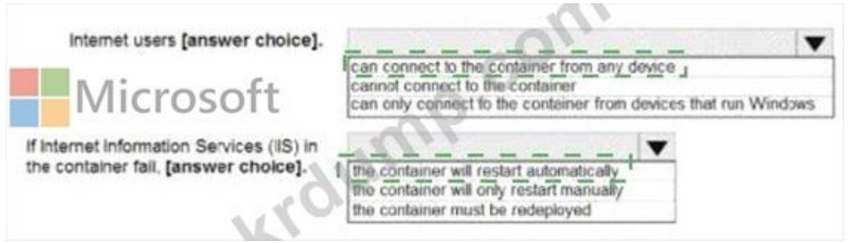

Explanation:

Box 1: can connect to the container from any device

In the policy "osType": "window" refer that it will create a container in a container group that runs Windows but it won't block access depending on device type.

Box 2: the container will restart automatically

Docker provides restart policies to control whether your containers start automatically when they exit, or when Docker restarts. Restart policies ensure that linked containers are started in the correct order. Docker recommends that you use restart policies, and avoid using process managers to start containers.

on-failure : Restart the container if it exits due to an error, which manifests as a non-zero exit code.

As the flag is mentioned as "on-failure" in the policy, so it will restart automatically

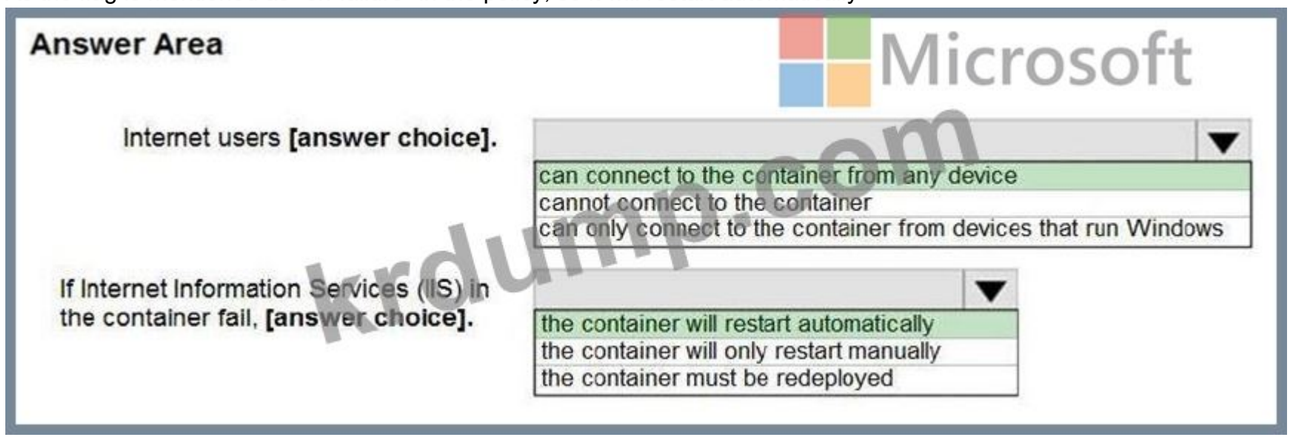

Reference:

https://docs.microsoft.com/en-us/cli/azure/container?view=azure-cli-latest

https://docs.docker.com/config/containers/start-containers-automatically/

NEW QUESTION: 37

□□ □□ □□□ □□□□ Azure Resource Manager □□□□ □□□□□□□. □ □□□□ 100□□ □□ □□□ □□□□ □ □□□□□.

□□□ □□□□□ □□□□□ □□□□ □□□□ □□□. □□□□□ □□ □□□□ □□□□ □□ □□□□ □□□.

□□□□□ □□□□□ □□□ □□□□ □□□?

A. Azure Active Directory(AD) ID □□ □ Azure □□

B. □□ □□□ □□ □ □□ □□

C. Azure Key Vault □ □□□ □□

D. Azure Storage □□ □ □□□ □□

Answer: (SHOW ANSWER)

According to the Microsoft Azure Administrator (AZ-104) study guide and Azure documentation, when deploying Azure resources through ARM templates, sensitive information such as administrative passwords must not be stored in plain text within the template or parameter files.
To securely store and reference such secrets, Azure Key Vault is the recommended service.
Azure Key Vault is a cloud service designed to store and manage secrets, encryption keys, and certificates securely. By integrating Key Vault with Azure Resource Manager templates, you can:
* Securely store passwords as secrets.
* Reference them dynamically in the template using a secure URI (reference function).
* Ensure the secret value is not exposed in deployment logs or ARM code.
* Control access using Access Policies or Azure RBAC.
The parameter section of the ARM template would reference the Key Vault as follows:
"adminPassword": {
"reference": {
"keyVault": {
"id": "/subscriptions/<subscription-id>/resourceGroups/<resource-group-name>/providers/Microsoft.KeyVault
/vaults/<keyvault-name>"
},
"secretName": "AdminPassword"
}
}
Thus, storing the administrative password in Azure Key Vault with a properly defined access policy is the Microsoft-verified, secure, and compliant approach.
Final Verified Answer: C. an Azure Key Vault and an access policy

NEW QUESTION: 38
Azure Resource Manager □□□□ □□□□ □□ □□ □□ □□□□ Windows Server 2022□ □□□□ □□ Azure □□ □□□ □□□ □□□□□.
□□ □□□ □□□ □□□ □□ □□ □□□□ NGINX□ □□□ □ □□□ □□□□ □□□.
□□□ □□□□ □□□?
A. Microsoft Intune □□ □□ □□□
B. Microsoft entra □□□□□□ □□□
C. Azure □□□ □□ □□□□ □□
D. Azure App Service□ □□ □□
Answer: C (LEAVE A REPLY)
When deploying Azure Virtual Machines (VMs) - individually or within a Virtual Machine Scale Set (VMSS) - administrators often need to automate post-deployment configuration tasks such as installing software, running configuration scripts, or setting up environment dependencies.
According to Microsoft Azure Administrator documentation, the Azure Custom Script Extension is designed precisely for this purpose. It downloads and executes scripts on Azure VMs after deployment. The extension integrates directly with Azure Resource Manager (ARM) templates, Azure CLI, or PowerShell, enabling automated installation and configuration of software packages like NGINX, IIS, or custom utilities.
The Custom Script Extension runs PowerShell scripts on Windows VMs and shell scripts on Linux VMs. In the case of deploying Windows Server 2022 VMs in a scale set, you can embed the Custom Script Extension JSON configuration into your ARM template, specifying the script's location in Azure Storage or GitHub and the command to execute (for example, Install-WindowsFeature Web-Server or an NGINX installer command).
This approach aligns with Microsoft's best practices for Azure automation and infrastructure-as-code, ensuring consistent, repeatable deployments across all instances in the scale set. It also minimizes manual intervention, reduces configuration drift, and ensures each VM instance has the required software components available immediately after provisioning.

NEW QUESTION: 39
Scope1□ □□ □□□ □□ □□□ □□□□□.

Scope1□ □□□ □□ □□□ □□□□□ □□□□ □□□.
Scope1□ □□□□ □□□ □□□□ □ □□□?
A. storage2□□ □□ □□□□□ □□
B. storage1 □ storage2□ □□□□ □ Blob
C. storage2□□ □□ □□□□, □□ □ □□ □□
D. storage1 □ storage2□ □□□□, Blob □ □□ □□
E. storage2□ □□□□, □□, □□ □□, □ □ □□□□ □□
Answer: E (LEAVE A REPLY)
In Microsoft Azure, encryption scopes are a StorageV2 (general-purpose v2) storage account feature that allows fine-grained control over encryption settings for data stored within a single account. According to Microsoft Azure Storage documentation, an encryption scope defines a specific encryption context that can be applied at the container or blob level and is supported in non-hierarchical namespace storage accounts (those without Data Lake Gen2 enabled).
In the given scenario:
* storage1 has Hierarchical namespace = Yes (Data Lake Storage Gen2 enabled).
* storage2 has Hierarchical namespace = No.
* The plan was to create an encryption scope named Scope1 in storage2.
* The technical requirement specifies that Scope1 must be used to encrypt storage services.
According to the Azure Administrator documentation on encryption scopes:
"Encryption scopes are supported for block blobs, append blobs, page blobs, Azure Files, queues, and tables in standard StorageV2 accounts. Encryption scopes are not supported in hierarchical namespace (Data Lake Gen2) enabled accounts." This means that Scope1-created in storage2, which does not have hierarchical namespace-can encrypt all blob data (containers and blobs) as well as file shares, queues, and tables.
However, storage1 cannot use encryption scopes because hierarchical namespace storage accounts (ADLS Gen2) manage encryption at the account level and do not support per-scope encryption.
Therefore, only storage2 can apply Scope1, and it can encrypt containers, blobs, file shares, queues, and tables.

NEW QUESTION: 40
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□ □ □□□□□. □ □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □ □ □□□ □□□ □□ □□, □□ □ □□□ □□□ □□ □□□ □□□□.
□ □□□ □□□ □□ □□□ □□ □□□ □ □□□□. □□□ □□□ □□□ □□ □□□ □□□□ □□□□.
image1□□□ □□□□ □□□□ Registry1□□□ Azure □□□□ □□□□□□ □□□□.
image1□ □□□□ □□□□ □□□□□ □□□□□ □□ □□ □□□□ □□□□□.
image1□ □□□□ □□□□ □□□□□ □□□ □ □□□ □□□.
□□ □□: Registry1□ □□ □□□ □□□□□ □□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
In Microsoft Azure, the Azure Container Registry (ACR) is a managed service that allows you to store and manage container images for all types of container deployments. When deploying an Azure Container Instance (ACI) from an ACR image, the deployment may fail if network configuration, authentication, or permissions are not correctly set.
The option "Use dedicated data endpoint" in ACR is designed to isolate network traffic for data operations (such as image push/pull) from the registry's control plane operations. However, enabling or disabling this feature does not affect authentication or deployment permissions to an Azure Container Instance.
According to the Azure Administrator Study Guide (Microsoft Official Documentation):

"To deploy a container instance from an Azure container registry, the registry must be accessible either publicly with proper authentication (admin user or service principal with AcrPull permission) or privately using a Virtual Network with Private Link (Premium tier). If you receive authentication or access errors, the solution is to verify credentials or network accessibility, not to enable a dedicated data endpoint." In this case, the root cause of the deployment error is most likely related to image access authentication or tier limitations, not the use of dedicated data endpoints. Therefore, selecting Use dedicated data endpoint will not resolve the deployment failure.
The verified solution as per Microsoft Learn and AZ-104 exam content is to either:
Enable the admin user or
Assign a managed identity or service principal with the AcrPull role to the container instance.
Hence, the proposed solution does not meet the goal.

NEW QUESTION: 41
□□ □□ □□□ Azure □□□□ □□□□□□ □□ Azure □□□ □□□□.

Name	Service tier
ContReg1	Premium
ContReg2	Standard
ContReg3	Basic

ACR □□□ □□□□ □□□□□ □□□ □□□ □□□□□.

Answer:

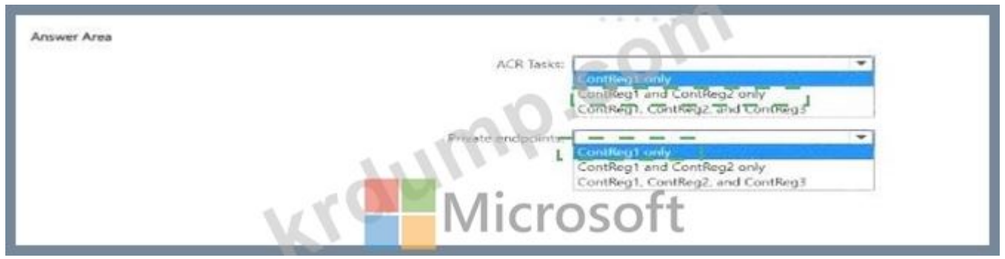

Explanation:
ACR Tasks: ContReg1 and ContReg2 only
Private endpoints: ContReg1 only
Azure Container Registry (ACR) provides multiple service tiers (SKUs)-Basic, Standard, and Premium- each offering different capabilities for scalability, performance, and features. These tiers determine what functionalities are available, including ACR Tasks and Private Link (Private Endpoints) integration.
Here's how each tier differs according to Microsoft Azure Administrator documentation (AZ-104 Exam Study Guide and Azure Docs):
Feature
Basic
Standard
Premium
ACR Tasks (Build automation, image building, and updating)
Supported

Supported
Supported
Private Link / Private Endpoints (Private access via Azure backbone)
Not Supported
Not Supported
Supported
Geo-replication
Not Supported
Not Supported
Supported
Content trust and RBAC integration
Supported
Supported
Supported
Analysis:
ACR Tasks:
ACR Tasks allow you to automate image builds and updates when base images or application code changes.
Supported in Standard and Premium tiers (as per Microsoft Docs: "ACR Tasks are available in Standard and Premium service tiers.").
Therefore, ContReg1 (Premium) and ContReg2 (Standard) support ACR Tasks.
ContReg3 (Basic) does not support advanced tasks automation at scale.
Answer: ContReg1 and ContReg2 only
Private Endpoints:
Private Link (Private Endpoints) enables private network connectivity to the ACR from a virtual network.
According to Azure documentation:
"Private Link for Azure Container Registry is supported only in the Premium service tier." This ensures that access to the registry occurs via Azure's backbone, preventing exposure to the public internet.
Answer: ContReg1 only (Premium)
Verified Microsoft Documentation Extract:
From Microsoft Learn: Azure Container Registry service tiers overview:
"The Premium tier provides advanced capabilities such as geo-replication, content trust, customer-managed keys, and Private Link support."
"ACR Tasks is available in both Standard and Premium tiers for continuous integration workflows and automated image building."

NEW QUESTION: 42
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□ □ □□□□□. □ □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □ □ □□□ □□□ □□ □ □□, □□ □□□□ □□□ □□ □ □□□□.
□ □□□ □□□ □□ □□□ □□ □□□ □ □□□□. □□□ □□□ □□□ □□ □□□ □□□□ □□□□.
image1□□□ □□□□ □□□□ Registry1□□□ Azure □□□□ □□□□□□ □□□□.
image1□ □□□□ □□□□ □□□□□ □□□□□ □□ □□ □□□□ □□□□□.
image1□ □□□□ □□□□ □□□□□ □□□ □ □□□ □□□.
□□□: Registry1□ □□ □□□ □□□ □□□□ □□□□□□.
□□□ □□□ □□□□□?
A. □

B. □□□

Answer: (SHOW ANSWER)

When deploying an Azure Container Instance (ACI) using an image from an Azure Container Registry (ACR), the deployment must authenticate to the registry to pull the image. By default, ACR is secured, and you cannot pull images anonymously unless they are public.

According to the Azure Administrator study guide, enabling the Admin user option in ACR generates two access credentials (username and password) that can be used for authentication when pulling container images. When the Admin user setting is disabled, you cannot authenticate directly using those credentials, and automated deployments may fail with authentication errors.

By enabling the Admin user setting, Azure creates a service-level user account for the registry with permissions to push and pull images. These credentials can then be referenced securely in your container instance configuration - either through the Azure portal, ARM template, or CLI parameters such as -- registry-username and --registry-password.

This configuration allows ACI to authenticate successfully against the ACR and pull the required image (image1) for deployment.

Therefore, enabling the Admin user for the container registry meets the goal of allowing the deployment of the container instance using image1.

NEW QUESTION: 43

□□ □□ □□□ □□□□ □□□ Azure AD□ □□□□□ □□□ □□□□.

Name	User type	On-premises sync enabled
User1	Member	No
User2	Member	Yes
User3	Guest	No

□□□□ JobTitle □ UsageLocation □□□ □□□□ □□□.

□□ □□□□ Azure AD □□□ □□□ □ □□□? □□ □□□□ □□□ □□□ □□□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer:

Explanation:

In a hybrid Azure Active Directory (Azure AD) environment, where user identities can be cloud-only or synchronized from on-premises Active Directory (AD), attribute management depends on the source of authority for each user account.

The source of authority determines where you can modify a user's attributes:

If a user is synced from on-premises Active Directory (on-premises sync enabled = Yes), certain attributes are read-only in Azure AD and must be edited in the on-premises AD. These attributes include:

Job title

Department

Office

Manager

Company name

If a user is cloud-only (not synchronized) or a guest user, you can modify those same attributes directly in Azure AD via the Azure portal, Microsoft Graph, or PowerShell.

According to Microsoft Learn (Azure AD attributes and synchronization rules):

"When Azure AD Connect synchronizes identities, certain user attributes such as job title, department, and manager are mastered in on-premises Active Directory and become read-only in Azure AD. However, attributes like usage location can always be edited in Azure AD, as they are required for license assignment." Applying to this question User Type On-premises sync enabled Editable in Azure AD User1 Member No Editable (Cloud-only) User2 Member Yes On-premises sync (Read-only for JobTitle) User3 Guest No Editable in Azure AD JobTitle User1 - Editable (cloud-only member) User2 - Not editable (synced from on-prem AD) User3 - Editable (guest accounts managed in Azure AD)

Answer: User1 and User3 only

UsageLocation

This attribute must be set in Azure AD for licensing and service access.

It can be edited for all users - even for synced users or guests - because Azure AD maintains it as a cloud- only attribute, not synchronized from on-prem AD.

Answer: User1, User2, and User3

Final Verified Answers (Based on Microsoft Azure Administrator Documentation):

Attribute

Editable Users

JobTitle

User1 and User3 only

UsageLocation
User1, User2, and User3
Microsoft Official Documentation Extract:
"Attributes such as job title, department, and office location are sourced from on-premises AD and cannot be modified in Azure AD for synchronized users. The usage location attribute, however, is always managed in Azure AD." (Source: Microsoft Learn - Azure AD Connect sync: Attributes synchronized to Azure AD and Manage user profile attributes in Azure AD)

NEW QUESTION: 44
storage1□□□ Azure Storage □□□ □□□□.
Azure □□□□ □□□□□□ □□□□ App1□ App2□□ Azure App Service □□ □□□□. □ □□ □□□ ID□ □□□□□.
App1□ App2□ storage1□□ BLOB□ □□ □ □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
* □□□□ □□□ □□ □□□□□□.
* App2□ □□ 30□ □□ storage1□□□ □□ □ □□□ □□□□□.
□ □□ □□ storage1□□ □□□ □□□□ □□□? □□□□ □□ □□□□ □□□ □□□ □□□□□. □□: □□ □□□ 1□□□□.

Answer:

Answer Area

Explanation:

The question involves two applications - App1 and App2 - that both need read access to blobs in an Azure Storage account (storage1). Both apps are running in Azure container instances and use managed identities for authentication.

Let's analyze the requirements and correct configuration for each app based on Azure's security and access control models.

App1 - Minimize Secrets

App1 uses a managed identity, meaning it can be authenticated to Azure services without any stored credentials or secrets.

The best practice is to assign Azure RBAC permissions (role-based access control) directly at the storage account or container level.

By using Access control (IAM), you can assign the Storage Blob Data Reader role to App1's managed identity.

This method uses Azure AD-based authentication, requires no SAS tokens or access keys, and minimizes secret management.

Access is continuous until the role is removed or modified.

Therefore, App1 # Access control (IAM)

App2 - Temporary 30-day Access

The requirement specifies that App2 should be able to read blobs only for 30 days.

Azure RBAC roles (IAM) do not provide time-bound permissions.

The appropriate way to grant time-limited access is through a Shared Access Signature (SAS).

A SAS token defines permissions, resource scope (e.g., container or blob), and an expiry time - making it ideal for temporary or limited access scenarios.

You can generate a SAS token valid for 30 days and assign it to App2.

Therefore, App2 # Shared access signatures (SAS)

Why Not Access Keys or Advanced Security

Access Keys: Grant full control (read/write/delete) to the storage account - not secure or granular, and they cannot be time-bound.

Advanced Security: Refers to configurations such as firewall rules or encryption; not directly related to granting app access.

Microsoft Azure Administrator Documentation Extract (AZ-104 Study Guide Reference):

"To enable secure access for applications, use Azure AD authentication with managed identities and assign appropriate RBAC roles via Access control (IAM). For temporary or limited access, use Shared Access Signatures (SAS) to specify permissions and expiry times." (Source: Microsoft Learn - Secure access to Azure Storage with Azure AD, SAS, and managed identities.)
Final Verified Answer:
App1: Access control (IAM)
App2: Shared access signatures (SAS)

NEW QUESTION: 45
5,000□□ □□□ □□□ □□□ Microsoft Entra □□□□ □□□□.
AdminUser1□□□ □□□ □□□ □□□ □□□ □□□□.
AdminUser1□ □□□ □□□ □□ □□□ □□□□ □□□.
□□□ □□ □□□□ □□□ □□ □□□?
A. □□ □□□□□□ □□□ □□□ □ □□□ □□□□□.
B. □□□□ □□ □□□□□□ □□□□ □□□ □□□□□.
C. □□□□ □□□□□□ □ □□□□□ □□□□□.
Answer: B (LEAVE A REPLY)
In Microsoft Entra ID (formerly Azure Active Directory), roles are assigned to users to delegate administrative permissions in a least-privilege manner. The User Administrator role allows a user to manage other users and groups - for example, creating and managing user accounts, resetting passwords for non- administrators, and managing user group memberships.
To assign a role such as User Administrator, you must use the Directory role blade within the user's account properties in the Azure portal.
Step-by-step according to Microsoft documentation:
* Sign in to the Azure Portal using an account that has one of the following roles:
* Global Administrator
* Privileged Role Administrator
* Navigate to Azure Active Directory # Users # select AdminUser1.
* Under Manage, select Directory role.This blade shows all current role assignments for the selected user.
* Click Add assignment (or Modify role).
* Select the User Administrator role from the list of available directory roles, then click Add.
Once this is completed, AdminUser1 will have administrative permissions limited to user management activities within the tenant.
Why other options are incorrect:
* A. From the Groups blade, invite the user account to a new group:Group membership does not grant directory-level administrative permissions. Roles must be assigned at the directory role level, not via groups (unless using role-assignable groups configured for PIM).
* C. From the Licenses blade, assign a new license:Licenses determine service usage (e.g., Microsoft 365, Intune) and do not provide administrative privileges in Entra ID.
Extract from Microsoft Azure Administrator Documentation (Official Guide):
"To assign a role to a user, in the Azure portal, select the user, then under Manage select Directory role, and choose the role you want to assign." (Source: Microsoft Learn - Assign roles to users in Azure Active Directory)

NEW QUESTION: 46
Azure □□ □□□ □□□□ Subnet1□□□ □□□□ □□□□. NSG1□□□ □□□□ □□ □□(NSG)□ □□ □□□ □□ Subnet1, NSG1□ □□□□□.
NSG1□ Subnet1□ □□□□ Azure Portal□ □□□□ □□□□ □□□ □□□□ □□□. □□□□ □□ ...□ □□□ □ □□□ □□□.
□□□□ Destination□ □□□ □□□□ □□□?
A. □□□ □□
B. IP □□
C. □□□□□□ □□ □□

D. ☐☐

Answer: A (LEAVE A REPLY)

This question tests your understanding of Network Security Groups (NSGs) and how to configure them to control outbound traffic - specifically to block access to the Azure portal while allowing other internet traffic.

Scenario Summary

You have a subnet (Subnet1) containing Azure virtual machines.

The subnet is associated with NSG1, which currently has only the default rules.

You must prevent access to the Azure portal while allowing other internet access.

The Azure portal is accessed through HTTPS (TCP port 443) at the following URL:

https://portal.azure.com

Azure provides a convenient way to manage outbound/inbound traffic for Microsoft services using Service Tags.

What are Service Tags?

A Service Tag is a predefined identifier for a specific Microsoft service or group of IP address prefixes managed by Azure.

Examples include AzureCloud, Storage, Sql, AppService, AzurePortal, etc.

When you use a service tag in a network security rule, Azure automatically manages the underlying IP ranges.

This ensures the rule stays up to date even if Microsoft changes IP addresses for that service.

Applying to the Scenario

To block traffic from your subnet to the Azure portal, you can create an Outbound NSG rule in NSG1 with:

Direction: Outbound

Action: Deny

Protocol: TCP

Port: 443 (HTTPS)

Destination: Service tag = AzurePortal

This ensures that:

All outbound connections to the Azure portal are blocked.

Other outbound HTTPS connections (e.g., general internet browsing, APIs, etc.) remain unaffected, since they do not use the AzurePortal tag.

Using a Service Tag is the correct and Microsoft-recommended method for this configuration because it is specific, automatically maintained, and minimal in administrative effort.

Why Not the Other Options

Option

Reason for Incorrectness

B). IP addresses

Not recommended because Azure portal IPs change frequently. Managing static IP lists would be error-prone and unscalable.

C). Application security group (ASG)

Used to group VM NICs for intra-subnet traffic control, not for targeting external Azure services.

D). Any

Would block all outbound traffic, not just Azure portal, violating the requirement to "connect to other internet destinations."

Verified Microsoft Learn Documentation Extract

"Service tags represent a group of IP address prefixes from specific Azure services. You can use service tags in NSG rules to allow or deny traffic for the corresponding Azure service. The tag 'AzurePortal' can be used to control access to the Azure portal." (Source: Microsoft Learn - Use service tags to define network access controls on NSG rules, AZ-104 Exam Guide)

Final Verified Answer: A. Service tag

Use the Service Tag AzurePortal in the NSG rule's Destination field to block outbound access to the Azure portal while allowing all other internet connections.

NEW QUESTION: 47
□□□ □□□□□ □□ □□□ □□□□□ □□□□ □□□?
A. □□
B. □□□ □□
C. □□ □□
D. □□ □□□
Answer: C (LEAVE A REPLY)
Cost analysis: Correct Option
In cost analysis blade of Azure, you can see all the detail for custom time span. You can use this to determine expenditure of last few day, weeks, and month. Below options are available in Cost analysis blade for filtering information by time span: last 7 days, last 30 days, and custom date range. Choosing the first option (last 7 days) auditors can view the costs by time span.
Cost analysis shows data for the current month by default. Use the date selector to switch to common date ranges quickly. Examples include the last seven days, the last month, the current year, or a custom date range. Pay-as-you-go subscriptions also include date ranges based on your billing period, which isn't bound to the calendar month, like the current billing period or last invoice. Use the <PREVIOUS and NEXT> links at the top of the menu to jump to the previous or next period, respectively. For example, <PREVIOUS will switch from the Last 7 days to 8-14 days ago or 15-21 days ago.

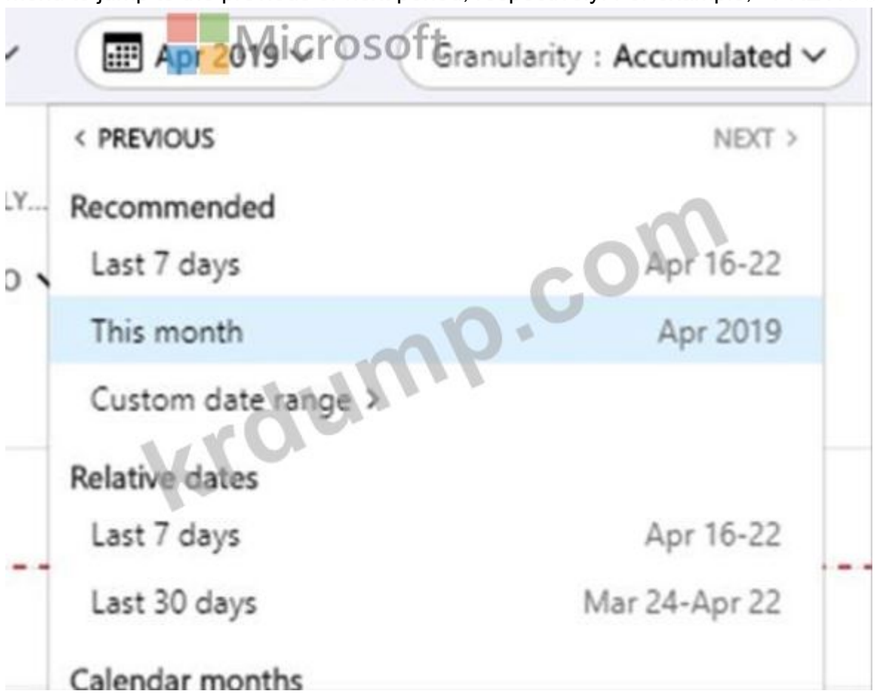

Invoice: Incorrect Option
Invoices can only be used for past billing periods not for current billing period, i.e. if your requirement is to know the last week's cost then that also not filled by invoices because Azure generates invoice at the end of the month. Even though Invoices have custom timespan, but when you put in dates for a week, the pane would be empty. Below is from Microsoft document:

Resource Provider: Incorrect Option

When deploying resources, you frequently need to retrieve information about the resource providers and types. For example, if you want to store keys and secrets, you work with the Microsoft.KeyVault resource provider. This resource provider offers a resource type called vaults for creating the key vault. This is not useful for reviewing all Azure costs from the past week which is required for audit.

Payment method: Incorrect Option

Payment methods is not useful for reviewing all Azure costs from the past week which is required for audit.

Reference:

https://docs.microsoft.com/en-us/azure/cost-management-billing/costs/quick-acm-cost-analysis

https://docs.microsoft.com/en-us/azure/cost-management-billing/manage/download-azure-invoice-daily- usage-date

NEW QUESTION: 48

□□ □□□ App Service □□□ □□ □□□□.

Name	Operating system	Location
ASP1	Windows	West US
ASP2	Windows	Central US
ASP3	Linux	West US

□□ □□ □□□ Azure □□□ □□ □□□□□.

Name	Runtime stack	Location
WebApp1	.NET Core 3.0	West US
WebApp2	ASP.NET 4.7	West US

□ □□ □□□ □ □□ App Service □□□ □□□□ □□□.

□□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

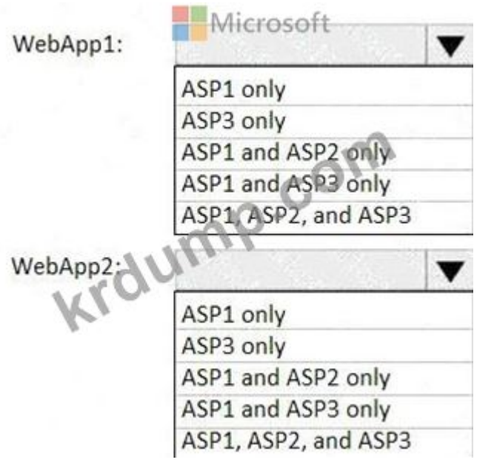

Answer:

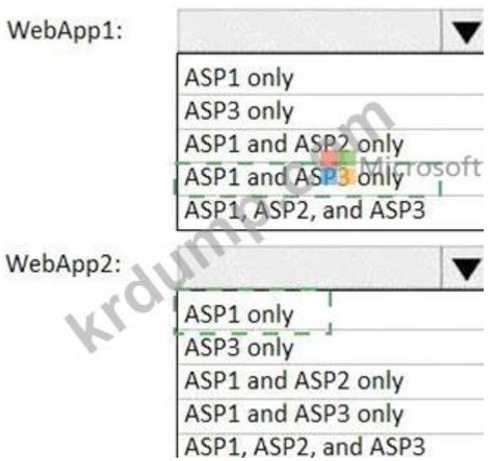

Explanation:

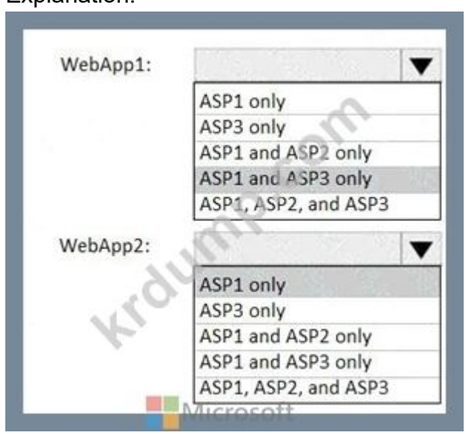

In Azure App Service, a Web App must be hosted in an App Service Plan (ASP) that meets three key compatibility conditions:

The App Service Plan and the Web App must be in the same region.
The operating system (Windows or Linux) of the App Service Plan must match the Web App type.
The runtime stack of the Web App must be supported on the App Service Plan's OS.
Let's apply these rules to the scenario:
App Service Plans:
Name
Operating System
Location
ASP1
Windows
West US
ASP2
Windows
Central US
ASP3
Linux
West US
Web Apps:
Name
Runtime Stack
Location
WebApp1
NET Core 3.0
West US
WebApp2
ASP.NET 4.7
West US
1## WebApp1 (.NET Core 3.0, West US)
Region requirement: Must use an App Service Plan in West US # # ASP1 and ASP3 qualify.
Runtime requirement: .NET Core 3.0 supports both Windows and Linux App Service Plans.
(Microsoft Learn: ".NET Core apps can run on both Windows and Linux App Service plans.")
Therefore, WebApp1 can use ASP1 (Windows, West US) or ASP3 (Linux, West US).
Answer: ASP1 and ASP3 only
2## WebApp2 (ASP.NET 4.7, West US)
Region requirement: Must be in West US # # ASP1 and ASP3 qualify.
Runtime requirement: ASP.NET 4.7 is a Windows-only framework; it cannot run on Linux App Service Plans.
(Microsoft Learn: "ASP.NET (non-Core) apps require a Windows-based App Service Plan.")
Therefore, only ASP1 (Windows, West US) is compatible.
Answer: ASP1 only
Final Verified Answers:
Web App

Compatible App Service Plans
WebApp1
ASP1 and ASP3 only
WebApp2
ASP1 only
Microsoft Documentation Extract (Azure App Service):
"App Service plans must be in the same region as the web app."
"Windows App Service plans host ASP.NET, .NET Core, and Node.js apps."
"Linux App Service plans host .NET Core, Node.js, Python, PHP, Java, and custom containers."
".NET Framework (ASP.NET 4.x) applications can only run on Windows-based App Service plans." Hence, the verified and Microsoft-official answer is:
WebApp1 # ASP1 and ASP3 only
WebApp2 # ASP1 only

NEW QUESTION: 49
□□□ □□□ □□□ Azure □□□ □□□□.
□□ □□□ □□□□ VM1□□□ □ □□ □□□ □□□□□.
* □□-□□ □□
* □□ □□□□: VNet1
* NIC □□□□ □□ □□: NSG1
□□□ □□□ □□□□ VM1 □□□□ □□□□□□ □□□.
□□ □□□ □□□□ □□□?
A. VM1□ □□ □□ □□
B. VM1□ □□ □□□
C. NSG1□ □□ NSG □□ □□
D. NSG1□ □□ □□ □□
Answer: C (LEAVE A REPLY)
Traffic Analytics analyzes the network security group (NSG) flow logs to provide insights into traffic flow in your Azure cloud1. NSG flow logs are a feature of Network Watcher that allows you to view information about ingress and egress IP traffic through an NSG2. To use Traffic Analytics, you need to enable NSG flow logs for the network security groups you want to monitor1.
Diagnostic settings for VM1 or NSG1 are not required for Traffic Analytics. Diagnostic settings are used to stream log data from an Azure resource to different destinations such as Log Analytics workspace, Event Hubs, or Storage account3. Insights for VM1 are also not required for Traffic Analytics. Insights are a feature of Azure Monitor that provide analysis of the performance and health of an Azure resource4.

NEW QUESTION: 50
contoso?02 3□□□ □□□ □□□ □□□ □□□ Azure □□□ □□□□. Contoso 2023 □□□ □□□□ □□ □□ □□ □□ □□□□ □□□□ □□□□.
Contoso 2023 □□□□ □□□ □□ □□□ □□ □□□□□.

□□ □□ □□□ □□□□ □□□ Microsoft Entra □□□□ □□□□.

Name	Shared access signature (SAS) token for contoso2023
User1	User delegation SAS with the maximum available permissions
User2	Service SAS with the maximum available permissions
User3	Account SAS with the maximum available permissions

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

No

No

No

In the given scenario, the storage account "contoso2023" has Allow storage account key access set to Disabled. This means that access keys and account-level Shared Access Signatures (SAS) cannot be used to authorize requests to the storage account. Only Azure AD authentication and user delegation SAS tokens are supported when key-based access is disabled.

Let's analyze each user according to Azure's documented storage access behavior:

* User1 - User delegation SAS:
* A user delegation SAS is based on Microsoft Entra (Azure AD) credentials and requires blob service (not file service).
* Since the user delegation SAS can only access Blob storage, and the question specifies cont1 (Blob container), this would normally allow access if Azure AD permissions exist. However, if no RBAC permissions are granted, access is denied by default.
* User2 - Service SAS:
* A service SAS uses a storage account key to sign the token.
* When Allow storage account key access is disabled, service SAS becomes invalid, because it relies on the account key.
* Therefore, User2 cannot access blob or file data using a service SAS in this configuration.
* User3 - Account SAS:
* An account SAS is also generated using the storage account key.
* With key access disabled, account SAS tokens are blocked, and any operation using them fails with an authorization error.
* Thus, User3 cannot access the content either in the blob container or the file share.

According to Microsoft Azure Administrator Documentation:

"When you disable storage account key access, requests that use the account access keys for authorization fail. Shared access signatures (SAS) that are signed with an account key or a service SAS become invalid. Only user delegation SAS tokens, which use Azure AD credentials, are supported." This aligns with the behavior described in Azure Storage security documentation and AZ-104 study guides.

NEW QUESTION: 51

□□ □□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

Name	Kind	Region
storage1	StorageV2	Central US
storage2	BlobStorage	West US
storage3	BlockBlobStorage	West US
storage4	FileStorage	East US

Appl□□□ □□□ □□ □□ Azure □□□ □□□□□.
Appl□ □□□□ □□□. □□□□ □□□ □□□□□ □□□.
□□ □□□□ □□ □□□□ □□□ □□□□ □□□?
A. □□□□1
B. □□□□2
C. □□□□3
D. □□□□4
Answer: D (LEAVE A REPLY)
To back up a web app, you need to configure a custom backup that specifies a storage account and a container as the target for the backup1. The storage account must be in the same subscription as the web app, and the container must be accessible by the web app2. The backup size is limited to 10 GB, and the backup frequency can be configured to minimize costs.
According to the table, storage1 is the only storage account that meets these requirements. Storage1 is in the same subscription and region as the web app, and it is a general-purpose v2 account that supports custom backups. Storage2 and storage3 are in a different region than the web app, which may incur additional costs for data transfer. Storage4 is a FilesStorage account, which does not support custom backups.
Therefore, you should use storage1 as the target for the backup of your web app. To configure a custom backup, you can follow these steps:
In your app management page in the Azure portal, in the left menu, select Backups.
At the top of the Backups page, select Configure custom backups.
In Storage account, select storage1. Do the same with Container.
Specify the backup frequency, retention period, and database settings as needed.
Click Configure.
At the top of the Backups page, select Backup Now.

NEW QUESTION: 52
□□□ □□ □□□ □□□□□ VM1□ VM2□ □□ □□□ □□□□ □□□.
□□ □ □□ □□□ □□□□ □□□□ □□□? □□□□ □□ □□□ □□ □□ □□□ □□ □□□□ □□□ □□□ □□□□ □□□□□.

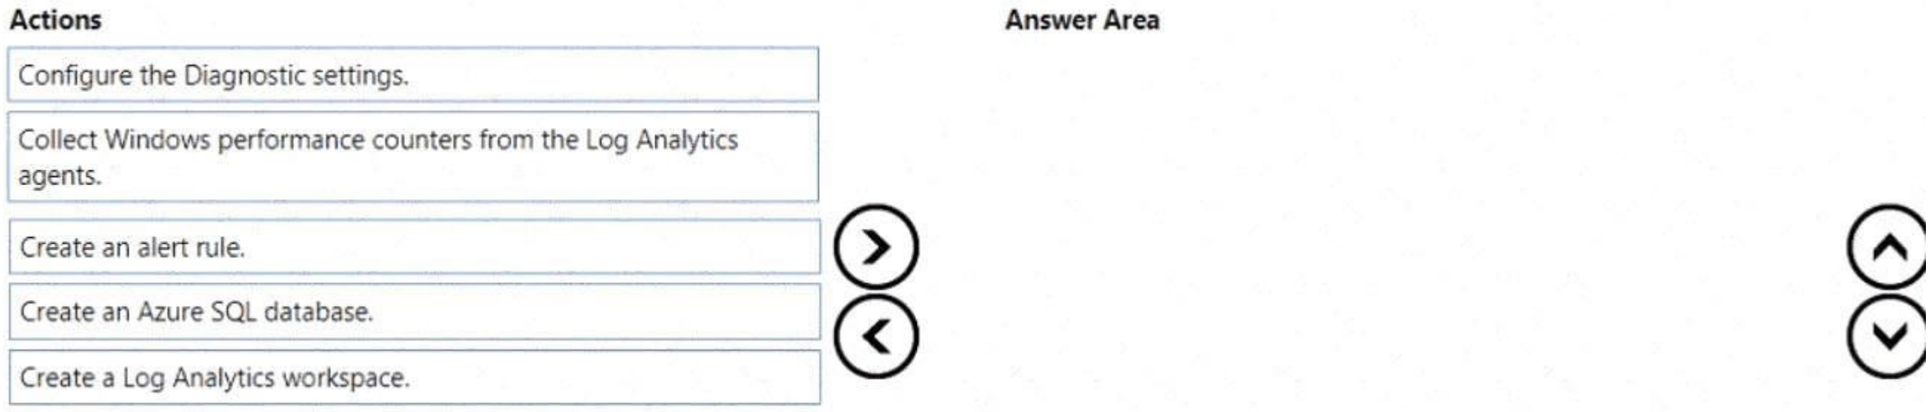

Answer:

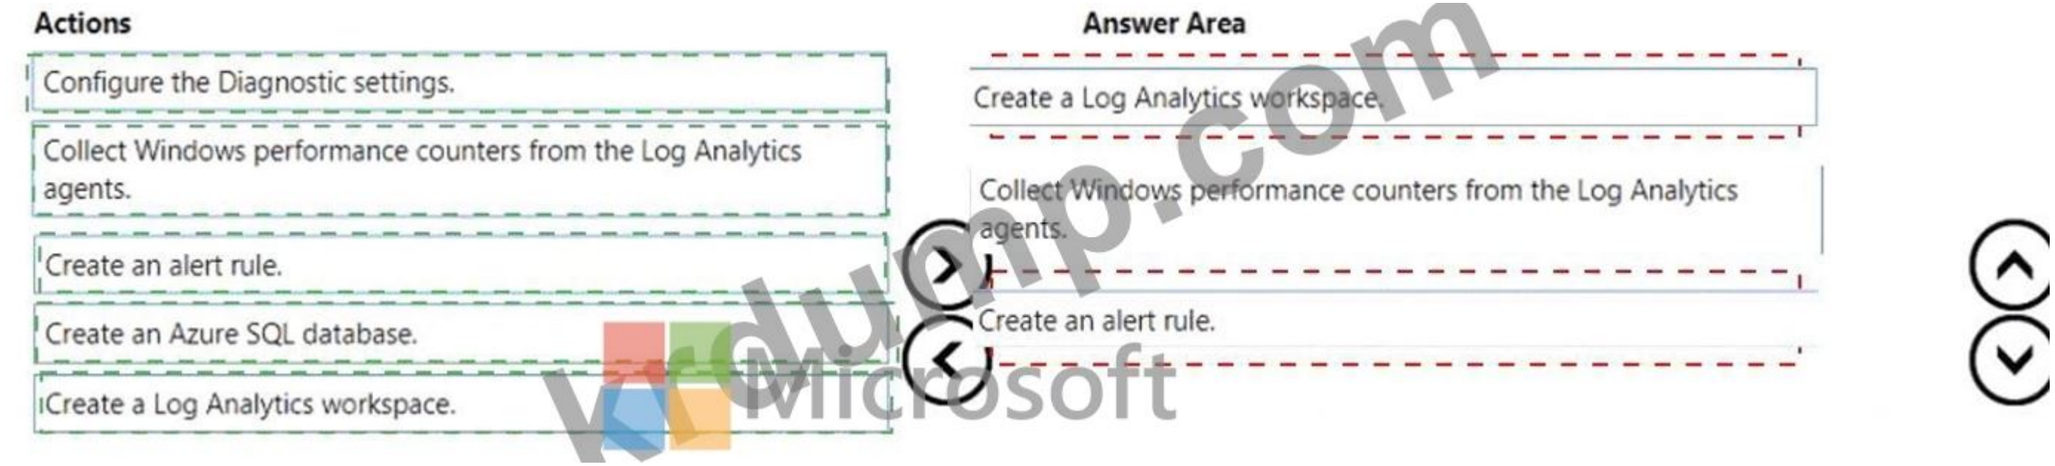

Explanation:

1## Create a Log Analytics workspace.

2## Collect Windows performance counters from the Log Analytics agents.

3## Create an alert rule.

The question states:

"You need to configure the alerts for VM1 and VM2 to meet the technical requirements." The technical requirement from the case study specifies:

"Trigger an alert if VM1 or VM2 has less than 20 GB of free space on volume C." This requirement involves monitoring performance metrics (disk space) and generating alerts based on those metrics. According to Microsoft Azure monitoring and management documentation, the process to configure such alerts involves using Azure Monitor and Log Analytics.

Step-by-Step Verified Solution:

Step 1: Create a Log Analytics workspace

A Log Analytics workspace is required to store and analyze logs and performance data collected from virtual machines.

Azure Monitor uses this workspace as the central repository for performance counters, events, and logs from connected agents.

From the Microsoft Documentation:

"Before you can collect and query data from virtual machines, you must create a Log Analytics workspace in your subscription." (Source: Azure Monitor and Log Analytics Guide) Step 2: Collect Windows performance counters from the Log Analytics agents After creating the workspace, you must connect VM1 and VM2 to the workspace and configure the Windows performance counters that you want to monitor.

In this case, you would collect the LogicalDisk(%) Free Space counter for C: drive.

Azure Monitor agent or Log Analytics agent collects these metrics and sends them to the workspace for analysis.

"To monitor system performance, configure the agent to collect performance counters such as available memory or free disk space." (Source: Azure Monitor Performance Counters Documentation) Step 3: Create an alert rule Once performance data is being collected, you can create an alert rule in Azure Monitor based on a Kusto query or metric threshold.

You would define a condition such as:

LogicalDisk | where FreeSpaceMB < 20480

and configure it to trigger an alert when free space on volume C drops below 20 GB.

This alert can notify via email, action group, or automation runbook.

"Alerts in Azure Monitor proactively notify you when important conditions are found in your monitoring data. Alerts can trigger automated actions or notifications." (Source: Azure Monitor Alerts Overview) Incorrect Options (Eliminated):

Configure the Diagnostic settings:

Used for collecting activity logs or resource logs, not performance counters from VMs.

Create an Azure SQL database:
Not relevant to the scenario of monitoring disk space.

NEW QUESTION: 53
□□ □□ □□□ □□ IP □□□ □□ Azure □□□ □□□□.

Name	IP version	SKU	Tier	IP address assignment
IP1	IPv4	Standard	Regional	Static
IP2	IPv4	Standard	Global	Static
IP3	IPv4	Basic	Regional	Dynamic
IP4	IPv4	Basic	Regional	Static
IP5	IPv6	Standard	Regional	Static

FW1□□□ Azure Firewall Premium □□□□□ □□□ □□□□□.
□□ IP □□□ □□□ □ □□□?
A. IP2 □□
B. IP1 □ IP2□ □□
C. IP1, IP2, IP5□
D. IP1, IP2, IP4, IP5□ □□
Answer: B (LEAVE A REPLY)
https://learn.microsoft.com/en-us/azure/virtual-network/ip-services/configure-public-ip-firewall Azure Firewall is a cloud-based network security service that protects your Azure Virtual Network resources. Azure Firewall requires at least one public static IP address to be configured. This IP or set of IPs are used as the external connection point to the firewall. Azure Firewall supports standard SKU public IP addresses. Basic SKU public IP address and public IP prefixes aren't supported.

NEW QUESTION: 54
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□ □ □□□□□. □ □□□ □□□ □□□ □□□ □ □□ □□□ □□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□, □□ □□ □□□□ □ □□ □□ □ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□ □□ □□□ □□ □□□ □□□□ □□□□.
Azure □□□ □□ □□□ □□□ □□□□□ □ □□□ □□□ Admin1□□□ Azure Active Directory(Azure AD) □□□□□ □□□□□□ □□□□ □□□.
□□ □□: □□ □□□□ Traffic Manager Contributor □□□ Admin1□ □□□□□.
A. □
B. □□□
Answer: (SHOW ANSWER)
The Traffic Manager Contributor role is not related to Traffic Analytics. Traffic Manager is a service that provides DNS-based load balancing and traffic routing across different regions and endpoints. Traffic Manager Contributor is a role that allows you to create and manage Traffic Manager profiles, endpoints, and geographies1.
Traffic Analytics is a service that provides visibility into user and application activity in your cloud networks.
Traffic Analytics analyzes Azure Network Watcher network security group (NSG) flow logs to provide insights into traffic flow in your Azure cloud. With Traffic Analytics, you can visualize network activity, identify hot spots, secure your network, optimize your network deployment, and pinpoint network misconfigurations2.
To enable Traffic Analytics for an Azure subscription, you need to have a role that grants you the following permissions at the subscription level:
Microsoft.Network/applicationGateways/read
Microsoft.Network/connections/read
Microsoft.Network/loadBalancers/read
Microsoft.Network/localNetworkGateways/read

Microsoft.Network/networkInterfaces/read
Microsoft.Network/networkSecurityGroups/read
Microsoft.Network/publicIPAddresses/read
Microsoft.Network/routeTables/read
Microsoft.Network/virtualNetworkGateways/read
Microsoft.Network/virtualNetworks/read
Microsoft.OperationalInsights/workspaces/*
Some of the built-in roles that have these permissions are Owner, Contributor, or Network Contributor3.
However, these roles also grant other permissions that may not be necessary or desirable for enabling Traffic Analytics. Therefore, the best practice is to use the principle of least privilege and create a custom role that only has the required permissions for enabling Traffic Analytics4.
Therefore, to meet the goal of ensuring that an Azure AD user named Admin1 is assigned the required role to enable Traffic Analytics for an Azure subscription, you should create a custom role with the required permissions and assign it to Admin1 at the subscription level.

NEW QUESTION: 55
Microsoft 365 □□□□ contoso.com□□□ Azure Active Directory(Azure AD) □□□□ □□□□.
User1, User2, User3□□□ □ □□ □□□□□ Library1□□□ □□ Microsoft SharePoint □□ □□□□□□ □□ □□□ □□□ □□□□□ □□□.
□□□ □□□ □□□□ □□□. □□□□ 180□ □ □□□ □□□□ □□□□□ □□ □□□.
□□ □ □□□ □□□□ □□□? □□□ □□ □□□ □□□□ □□□□□.
□□: □□ □□□ 1□□□□.
A. □□□ □□□ □□□ □□□□ □□ □□
B. □□□ □□□ □□□ □□□□ Office 365 □□
C. □□ □□□ □□□ □□□ □□□□ Office 365 □□
D. □□ □□□ □□□ □□□ □□□□ □□ □□
E. □□ □□ □□□ □□□ □□□□ □□ □□
Answer: B,C (LEAVE A REPLY)
You can set expiration policy only for Office 365 groups in Azure Active Directory (Azure AD).
Note: With the increase in usage of Office 365 Groups, administrators and users need a way to clean up unused groups. Expiration policies can help remove inactive groups from the system and make things cleaner.
When a group expires, all of its associated services (the mailbox, Planner, SharePoint site, etc.) are also deleted.
You can set up a rule for dynamic membership on security groups or Office 365 groups.

NEW QUESTION: 56
RG1□□□ □□□ □□□ □□□ Azure □□□ □□□□.
template1□□□ Azure Resource Manager(ARM) □□□□ □□□□ □□□□ □□□ □□□□□. □□□□ □□ □□ □□□ □□□□ □□□.
* RG1□ □□□ □□□□ □□□□□.
* □□□ □□□□ □□□□ □□ RG1□□ □□ □□□□ □□ □□□□□.
□□□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-resourcegroupname

Specifies the name of the resource group to deploy.

https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-mode

Specifies the deployment mode. The acceptable values for this parameter are:

-Complete: In complete mode, Resource Manager deletes resources that exist in the resource group but are not specified in the template.

- Incremental: In incremental mode, Resource Manager leaves unchanged resources that exist in the resource group but are not specified in the template.

NEW QUESTION: 57

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type
LB1	Load balancer
VM1	Virtual machine
VM2	Virtual machine

LB1□ □□ □□ □□ □□□□□.

Name	Type	Value
bepool1	Backend pool	VM1, VM2
LoadBalancerFrontEnd	Frontend IP configuration	Public IP address
hprobe1	Health probe	Protocol: TCP Port: 80 Interval: 5 seconds Unhealthy threshold: 2
rule1	Load balancing rule	IP version: IPv4 Frontend IP address: LoadBalancerFrontEnd Port: 80 Backend Port: 80 Backend pool: bepool1 Health probe: hprobe1

□□ □□ □□□ □□□□ □□□ □□□□ NAT □□□ □□ □□□□□.
□□ 3389□ □□□□ □□□□□ VM2□ □□ □□ □□□□ □□□□ □□□□□.
A. □□□□□ IP □□
B. □□ □□□
C. □□ □□ □□
D. □□□ □
Answer: A (LEAVE A REPLY)
To create an inbound NAT rule, you need to specify a frontend IP address and a frontend port for the load balancer to receive the traffic, and a backend IP address and a backend port for the load balancer to forward the traffic to1. According to the first table, LB1 has only one frontend IP address, which is
40.121.183.105. However, this frontend IP address is already used by the existing inbound NAT rule named rule1, which forwards port 80 to VM1 on port 802. Therefore, you cannot use the same frontend IP address and port for another inbound NAT rule.
To solve this problem, you need to create a new frontend IP address for LB1 before you can create the new inbound NAT rules. You can do this by using the Azure portal, PowerShell, or CLI3. After you create a new frontend IP address, you can use it to create the new inbound NAT rules that meet your requirements.

NEW QUESTION: 58
contoso.com□□□ Azure Active Directory(Azure AD) □□□□ □□□□.
500□□ □□ □□□□ □□□ □□□ □□□ □□□ CSV □□□ □□□□.
500□□ □□ □□□ □□□ □□ contoso.com□□ □□□ □□□ □□□ □□□□ □□□.
□□ □□: Azure Portal□ Azure AD□□ □□ □□□ □□□ □□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
In Microsoft Azure Active Directory (Azure AD), there is a clear distinction between internal users (members) and external users (guests). When you need to add a large number of external users (B2B collaborators) - such as those listed in a CSV file - you must use Azure AD B2B invitation processes, not the Bulk create user operation.
According to the Microsoft Azure Administrator Study Guide and Azure AD documentation, the Bulk create user operation in the Azure portal is designed only for internal user accounts within the organization's directory. It cannot be used to create guest user accounts. Guest users must be invited to the directory using either:
* The New-AzureADMSInvitation PowerShell cmdlet, which sends invitations to external users' email addresses and creates guest accounts (userType = "Guest").
* The Azure AD portal "Bulk invite" feature, which allows uploading a CSV file containing email addresses of external users to automate guest account creation.
The documentation explicitly states:
"To invite external users (B2B collaboration users) in bulk, you must use the Bulk invite feature or PowerShell with the New-AzureADMSInvitation cmdlet. The Bulk create operation is only supported for member users." (Source: Microsoft Learn - Azure Active Directory B2B collaboration and bulk operations guide) Therefore, since the scenario uses Bulk create user, it does not meet the goal of creating guest accounts for external users.

NEW QUESTION: 59
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.
VM1□□□ Azure □□ □□□ □□□□. VM1□ ARM1.json□□□ □□□ □□ Azure Resource Manager □□□□ □□□□ □□□□□□□.
VM1□ □□ □□□ □□□ □□ □□□□ □□□ □□□□.
VM1□ □□ □□ □□□□ □□□□ □□□.

□□ □□: □□□ □□□□□□ □□□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: A (LEAVE A REPLY)
When Azure schedules maintenance for a virtual machine, you can proactively move it to a new physical host by performing a self-service redeploy.
The Redeploy feature in the Azure portal allows you to:
* Move the VM to a new host node.
* Keep the same network interface, disks, and configuration.
* Resolve underlying host-level or platform issues proactively.
This action satisfies the requirement to move VM1 to a different host immediately and minimizes downtime.
This is explicitly documented in the Microsoft Learn - Redeploy Windows virtual machine to new Azure node guide.
Final Verified Answer: A. Yes

NEW QUESTION: 60
Azure □□ □□□ □□□□ □ □□□□ Azure Resource Manager□ □□□□.
Template1□□ □□ □□□□ □□□□ □□□□.

```
"location": {
    "type": "String",
    "defaultValue": "eastus",
    "allowedValues": [
        "canadacentral",
        "eastus",
        "westeurope",
        "westus" ]
}
```

Template1□ □□ □□□□ □□ □□□□ □□□□ □□□□.
"□□": "□□□"
Template1□ □□□ □□□□ □□ □□□□ □□□□ □□□□.

```
"type": "Microsoft.Compute/virtualMachines",
"apiVersion": "2018-10-01",
"name": "[variables('vmName')]",
"location": "westeurope",
```

Template1□ □□□□ □□ □□□ □□ □□ □□□ □□□□ □□□.
□□□ □□ □□□?
A. □□□ □□□ □□□ westus□ □□□□□.
B. □□ □□ □□ □□□ □□□□□
C. □□ □□□ □□□ westus□ □□□□□.
Answer: A (LEAVE A REPLY)

You can change the location in resources. Parameters used to define the value of some variables to be able to use in different places in the template resources. Resources are used only for complicated expressions. In any case, RM will only deploy from resources. In case the value is not mentioned directly, then it will check parameters if it is specified in the resources. Based on this question, the value of location is defined directly in resources. so you change the resources location value.
Use location parameter. To allow flexibility when deploying your template, use a parameter to specify the location for resources. Set the default value of the parameter to resourceGroup().location.
Reference:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/resource-location?tabs=azure- powershell
https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/template-syntax#resources

NEW QUESTION: 61
□□ □□□ □□ □□□ VM1□□□ Azure □□ □□□ □□ □□□□□.
VM1□ □□ □□□ □□□ □□□ □□ □□ □□ □□□□.

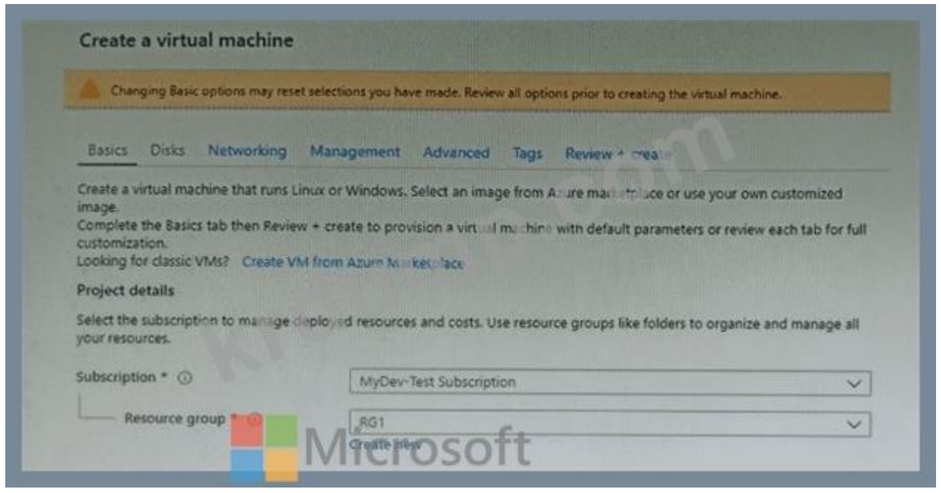

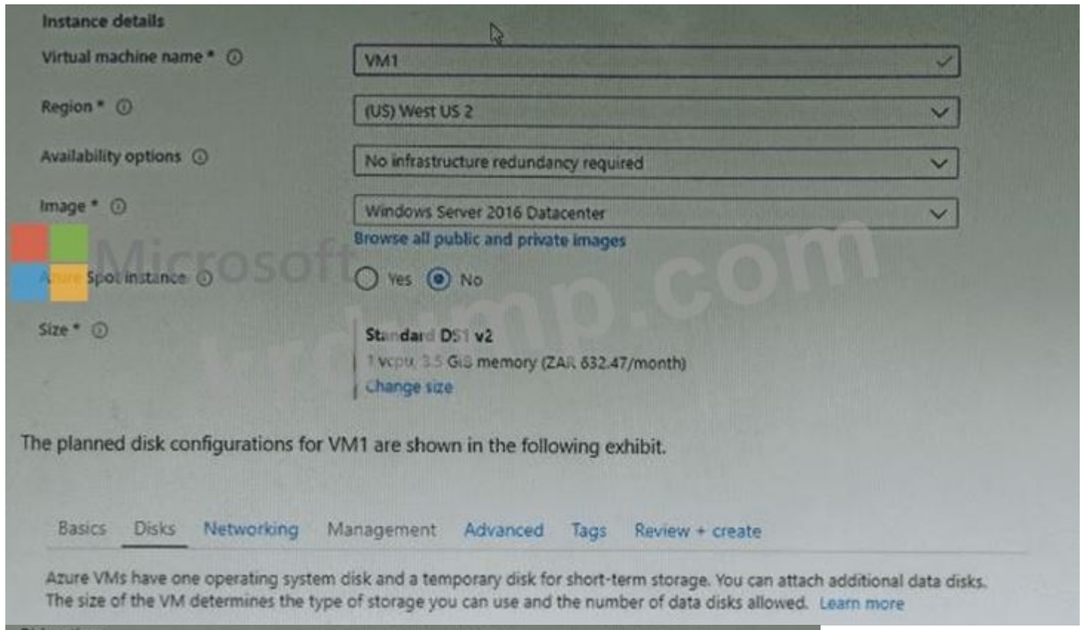

The planned disk configurations for VM1 are shown in the following exhibit.

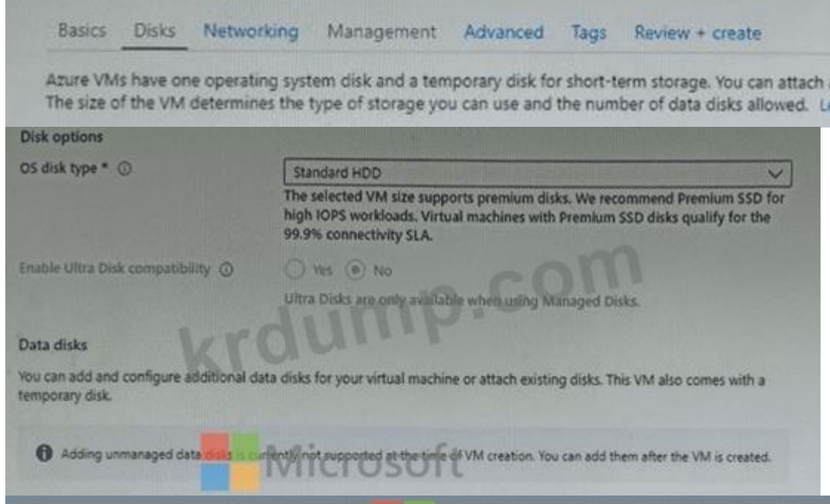

VM1□ □□□ □□□ □□□ □ □□□ □□□□ □□□.

□□ □ □□ □□□ □□□□ □□□? □ □□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

A. □□ □□□ □□

B. □□□ □□

C. OS □□□ □□

D. □□
E. □□□
Answer: A,B (LEAVE A REPLY)
https://docs.microsoft.com/en-us/azure/site-recovery/move-azure-vms-avset-azone https://docs.microsoft.com
/en-us/azure/virtual-machines/windows/create-portal-availability-zone https://docs.microsoft.com/en-us/azure
/virtual-machines/manage-availability https://docs.microsoft.com/en-us/azure/availability-zones/az- overview#availability-zones

NEW QUESTION: 62
VM1□□□ Azure □□ □□□ □□□□. VM1□ ARMljson□□□ □□□ □□ Azure Resource Manager □□□□ □□□□ □□□□□□□.
VM1□ □□ □□□ □□□ □□□□ □□□ □□□□□.
VM1□ □□ □□ □□□□ □□□□ □□□.
□□ □□: □□□ □□ □□□□□□ VM1□ □□ □□□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
According to the Microsoft Azure Administrator (AZ-104) official documentation and study guides, when a virtual machine (VM) in Azure is scheduled for maintenance, you can proactively move the VM to another host to minimize downtime. This process is called a self-service maintenance move.
When Azure schedules maintenance events that could affect the underlying host hardware, customers have the option to move their virtual machines within the same region and subscription to a healthy host by using either the Azure portal, PowerShell, or the Azure CLI. The operation that performs this action is called Redeploy or Maintenance Redeploy.
In this scenario, the proposed action - moving the VM to another subscription from the resource group blade
- does not satisfy the requirement. Moving a VM between subscriptions involves deallocation and redeployment, which does not guarantee that the VM will move to a new host to avoid maintenance. In fact, the VM will remain unavailable during the move process and may need to be reconfigured.
The correct solution to move a VM to another host immediately is to perform a self-service maintenance operation or redeploy the VM. This can be done by selecting the VM in the Azure portal, navigating to Help + support # Maintenance, and choosing "Move to a different host", or by using the Redeploy button under the VM settings.
Thus, moving the VM to another subscription does not meet the maintenance mitigation goal.

NEW QUESTION: 63
Sub1□□□ Azure □□□ □□□□.
□□ □□ □□□ □□□ □□□□ □□□ □□□□□□□ □□□ □□□□□.

Tier	Accessible from the Internet	Number of virtual machines
Front-end web server	Yes	10
Business logic	No	100
Microsoft SQL Server database	No	5

□□ □□ □□□ □□□□ □□□□ □□□□ □□□□ □□□.

* □ □□□ □□□□ □□ □□ □□ □□□ □□ □□ □□□ □□□□ □□□□□ □□□.

* □ □□□ SQL □□ □□□□□□ □□□□□.

□ □□ □□□ □□ Azure □□□□ □□□ □□□. □□□□□ □□ □□□□ □□□ □□□ □□□□□. □□: □□□ 1□□□□.

Answer Area

Ensure that communication between the web servers and the business logic tier spreads equally across the virtual machines:

an internal load balancer
- an application gateway that uses the Standard tier
- an application gateway that uses the WAF tier
- an internal load balancer
- a network security group (NSG)
- a public load balancer

Protect the web servers from SQL injection attacks:

an application gateway that uses the WAF tier
- an application gateway that uses the Standard tier
- an application gateway that uses the WAF tier
- an internal load balancer
- a network security group (NSG)
- a public load balancer

Answer:

Answer Area

Ensure that communication between the web servers and the business logic tier spreads equally across the virtual machines: an internal load balancer

- an application gateway that uses the Standard tier
- an application gateway that uses the WAF tier
- an internal load balancer
- a network security group (NSG)
- a public load balancer

Protect the web servers from SQL injection attacks: an application gateway that uses the WAF tier

- an application gateway that uses the Standard tier
- an application gateway that uses the WAF tier
- an internal load balancer
- a network security group (NSG)
- a public load balancer

Explanation:

Box 1: an internal load balancer

Azure Internal Load Balancer (ILB) provides network load balancing between virtual machines that reside inside a cloud service or a virtual network with a regional scope.

Box 2: an application gateway that uses the WAF tier

Azure Web Application Firewall (WAF) on Azure Application Gateway provides centralized protection of your web applications from common exploits and vulnerabilities. Web applications are increasingly targeted by malicious attacks that exploit commonly known vulnerabilities. Application gateway which uses WAF tier.

NEW QUESTION: 64

VM1□□□ Azure □□ □□□ Vault1□□□ Azure □ □□ □□ □□□ □□□□.

VM1□□ □ □□□ □(KEK)□ □□□□□ Azure Disk Encryption□ □□□ □□□□□. Azure Disk Encryption□ □□ Vault!□ □□□□ □□□.

Vault1□□ □□ □ □□ □□□ □□□□ □□□? □ □□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

A. □□□ □□ □□□□□.

B. □□□ □□ Azure Virtual Machines□ □□□□□.
C. □ □□ □□□ □□□□□.
D. □□□ □□□ □□□□.
E. □□ □□□□ □□ Azure Disk Encryption□ □□□□□.
Answer: A,C (LEAVE A REPLY)
To prepare Vault1 for Azure Disk Encryption, you need to perform the following actions on Vault1:
Create a new key. A key encryption key (KEK) is an encryption key that is used to encrypt the encryption secrets before they are stored in the key vault. You can create a new KEK by using the Azure CLI, the Azure PowerShell, or the Azure portal1. You can also import an existing KEK from another source, such as a hardware security module (HSM)2. The KEK must be a 2048-bit RSA key or a 256-bit AES key3.
Select Azure Disk Encryption for volume encryption. This is an advanced access policy setting that enables Azure Disk Encryption to access the keys and secrets in the key vault. You can select this setting by using the Azure CLI, the Azure PowerShell, or the Azure portal4. You must also enable access to Microsoft Trusted Services if you have enabled the firewall on the key vault.

NEW QUESTION: 65
□□ □□□□ □□□□□ 5,000□□ Blob□ □□□ Azure Storage □□□ □□□□.
□□□□ Blob □□□ □□□ □□□□ □□ Blob□ □ □ □□□ □□ □□□.
□□□□ □□□ □□□□ □□□?
A. JIT(Just-In-Time) VM □□□
B. □□ □□□ □□(SAS)
C. □□□ □□□ □□
D. □□ □□ □□
Answer: D (LEAVE A REPLY)
To ensure that users can view only specific blobs based on blob index tags, the correct solution is to configure Azure Role-Based Access Control (RBAC) role assignment conditions that use Azure Blob Index Tags for fine-grained data access.
According to the Microsoft Azure Administrator documentation, Blob index tags provide the ability to categorize and filter blobs within a storage account using key-value pairs. However, controlling access to blobs based on their index tags requires role assignment conditions, which is a feature of Azure RBAC conditional access for data actions.
This functionality is part of Azure Attribute-Based Access Control (ABAC). ABAC extends RBAC by adding conditions to role assignments so that access decisions can include resource attributes such as blob index tags or security principal attributes.
When you apply an RBAC role assignment with a condition, you can restrict access at the object level within a resource type. For example, you can grant a user the Storage Blob Data Reader role but restrict their access only to blobs that contain a specific tag, such as:
@Resource[Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/Department] StringEquals
'Finance'
This configuration ensures that even if multiple users have the same role, they can only view blobs whose tags match the condition defined in their role assignment.
Alternative options explained:
* A. Just-in-time (JIT) VM access # Used for securing virtual machines via Azure Security Center, not for blob data access control.
* B. Shared access signature (SAS) # Provides temporary access to storage resources but cannot filter access based on blob index tags.
* C. Stored access policy # Used to manage SAS tokens collectively but still does not support tag-based conditional access.
Therefore, the only mechanism that supports tag-based, condition-level access control for Azure Storage blobs is Role Assignment Conditions (Azure RBAC ABAC).
This aligns with the Microsoft Azure documentation on ABAC for Azure Storage, which explicitly states:
"You can use role assignment conditions in Azure Storage to control access to blobs based on blob index tags."

NEW QUESTION: 66
storage1□□□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

storage1□ □□ □□ □□□ □□(RABC) □□□ □□□ □ □□□ □□□ □□□□□. □□ storage1 □□□□ □□□ □□□ □ □□□ □□□□□?

A. □□□□□

B. □□ □□□ □□

C. □□□□

D. □□□□

E. □□□□ □ □□

F. □□ □□ □ □□□□

Answer: E (LEAVE A REPLY)

Azure role-based access control (RBAC) now supports role assignment conditions for finer-grained access management. Conditions are written using the Azure Resource Manager (ARM) condition language and allow you to enforce specific rules (for example, limit access to particular blobs or queues).

However, conditional access in RBAC is currently available only for data actions in Azure Storage accounts and Azure Key Vault. According to the Microsoft Learn documentation for Azure Storage RBAC with conditions, the following services support conditional role assignments:

* Blob storage (containers and blobs)

* Queue storage

This means that you can apply conditions on containers (for blobs) and queues, but not on file shares or tables.

Conditions can restrict access to:

* Specific container names or blob prefixes.

* Specific queue names or messages.

For example, you could allow a user to read blobs only under a given prefix or queue, enhancing least- privilege control.

Supported: Containers (Blob storage), Queues

Not supported: File shares, Tables

Microsoft Azure Reference (Conceptual Summary):

"You can add conditions to Azure role assignments for blob and queue data actions. Conditions are not yet supported for Azure Files or Tables." (Source: Microsoft Learn - Azure role assignment conditions for storage data actions)

NEW QUESTION: 67

RSV1□□□ Recovery Services □□ □□ □□□ □□□□. RSV1□□ 5□ □□ □□ □□□□ □□□□ 14□ □□ □□ □□□ □□□□ □□ □□□ □□□□.

RSV1□ VM1□ □□ □□□ □□□□□. VM1□ 8□ □□ □□□□□ □□ □ □□□□ □□□□□□.

VM1□ 8□ □ □□□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□□ □□□.

□□ □□□ □□ □□□?

A. VM1 □□□ □□□□□.

B. □□ □□ □□ □□ □□□ □□□□ VM1□ □□□□□.

C. VM1□ □□□□□.

D. □ □□ □□ □□□ □□□ □□□□ VM1□ □□□□□.

Answer: D (LEAVE A REPLY)

Azure Backup uses Recovery Services vaults (RSVs) to manage backup and restore operations for virtual machines. Each backup consists of:

Instant Restore Snapshots - retained for quick recovery (up to 5 days in this case).

Daily Recovery Points - stored in the Recovery Services vault based on the retention policy (14 days here).

In the scenario, VM1's website was updated eight days ago, meaning the restore point required is eight days old - beyond the 5-day instant snapshot retention period but within the 14-day daily backup retention window.

According to the Azure Backup documentation, restoring from a vault-stored daily recovery point (not instant snapshot) involves creating a new virtual machine because it ensures a non-disruptive restore and minimal downtime.

Here's how it works:
The backup is used to create a new VM in the same or different resource group.
Once the new VM is confirmed operational, you can redirect traffic or replace the original VM during a controlled maintenance window.
The "Replace existing" option, on the other hand, overwrites the current VM - leading to downtime and risk if validation fails.
Therefore, to restore VM1 to its state from eight days ago while minimizing downtime, you must first restore using the "Create new restore configuration" option, validate the restore, and then switch over seamlessly.

NEW QUESTION: 68
VM1□□□ Azure □□ □□□ □□□ Azure □□□ □□□□. VM1□ □□ □□ □□□□□ □□□□ □□ App1□□□ □□ □□ □□ □□□□□.
□□ □, App1□ □□□ □ VM1□ CPU □□□□ □□□□ □□□□.
□□ □□ VM1□ □□□□ □□□ □□□□ □□□ Runbook□ □□□□ □□□.
Runbook□ □□ □□□ □□□□ □□□?
A. Azure Performance Diagnostics □□□□□ VM1□ □□□□□.
B. VM1□ VM □□ □□□ □□□□□.
C. □□ □□□ VM1□ □□□□□.
D. □□□ □□ vCPU □□□□ □□□□.
E. VM1□ DSC(Desired State Configuration) □□□ □□□□□.
Answer: (SHOW ANSWER)
To create a scheduled runbook to increase the processor performance of VM1 at the end of each month, you need to modify the VM size property of VM1. This will allow you to scale up the VM to a larger size that has more CPU cores and memory. You can use Azure Automation to create a PowerShell runbook that changes the VM size using the Set-AzVM cmdlet. You can then schedule the runbook to run at the end of each month using the Azure portal or Azure PowerShell. For more information, see How to resize a virtual machine in Azure using Azure Automation1.

NEW QUESTION: 69
adatum.com□□□ Azure Active Directory(Azure AD) □□□□ □□□□. Adatum.com□□ □□ □□ □□ □□ □□□ □□□□ □□□□.

Name	Group type	Membership type	Membership rule
Group1	Security	Dynamic user	(user.city -startsWith "m")
Group2	Microsoft Office 365	Dynamic user	(user.department -notIn ["HR"])
Group3	Microsoft Office 365	Assigned	*Not applicable*

□□ □□ □□□ □□ □□□ □ □□ □□□ □□□ □□□□.

Name	City	Department	Office 365 license assigned
User1	Montreal	Human resources	Yes
User2	Melbourne	Marketing	No

User1□ User2□ □□ □□□ □□□□? □□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

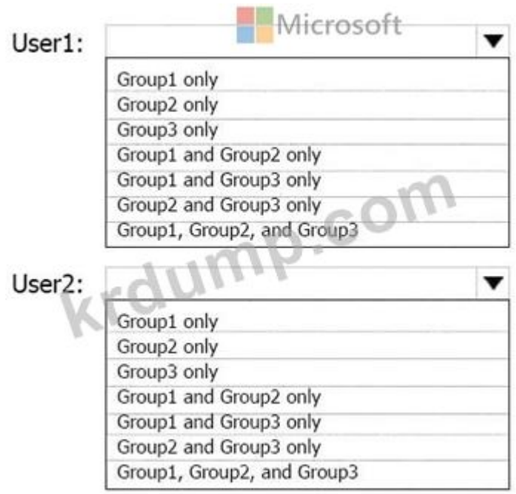

Answer:

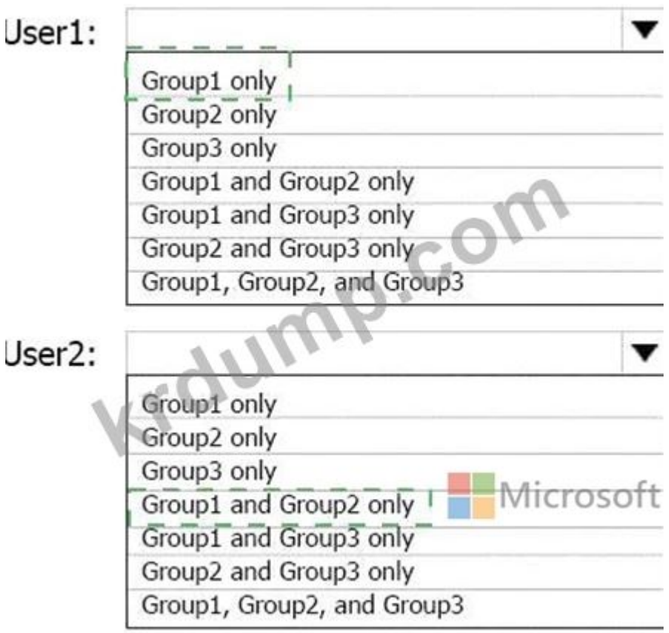

Explanation:

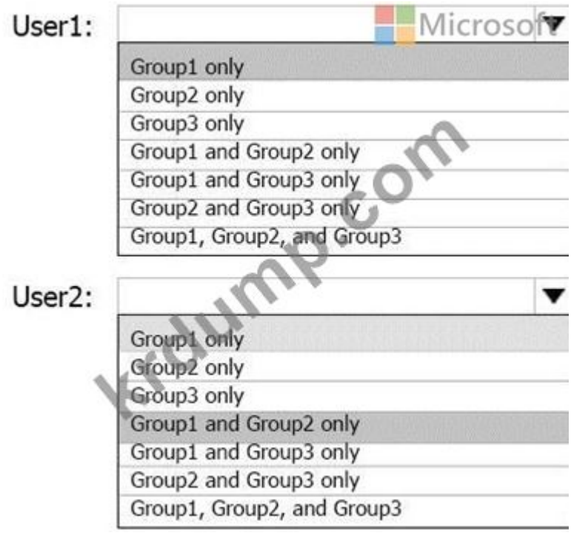

Azure Active Directory (Azure AD) supports three types of groups:
Security groups (for access and permissions),
Microsoft 365 groups (for collaboration tools such as Teams, SharePoint, and Outlook), Dynamic groups (with automatically managed memberships based on user attributes).
A dynamic user group uses membership rules to automatically add or remove users based on user properties (e.g., city, department, job title, etc.). These rules use Azure AD rule syntax, and they are case-insensitive.
Given Groups
Name
Group type
Membership type
Membership rule
Group1
Security
Dynamic user
(user.city -startsWith "m")
Group2
Microsoft 365
Dynamic user
(user.department -notIn ["HR"])
Group3
Microsoft 365
Assigned
Not applicable
Given Users
Name
City
Department
Office 365 license assigned

User1
Montreal
Human resources
Yes
User2
Melbourne
Marketing
No
Evaluate Group Membership
Group1 Rule:
(user.city -startsWith "m")
The condition includes any user whose city starts with 'm' (case-insensitive).
User1: City = Montreal # # Matches
User2: City = Melbourne # # Matches
Both User1 and User2 belong to Group1
Group2 Rule:
(user.department -notIn ["HR"])
The condition excludes users in the "HR" department.
Note: "Human resources" # "HR" (string comparison must match exactly).
User1: Department = Human resources # # Not "HR", so matches the rule.
User2: Department = Marketing # # Not "HR", matches the rule.
Both User1 and User2 belong to Group2
Group3:
Assigned group - membership must be manually configured.
No users are mentioned as being assigned # # Neither User1 nor User2 belongs.
However, here's a key Azure detail:
Dynamic Microsoft 365 groups (like Group2) require that users have a valid Microsoft 365 license. Users without a license (User2 in this case) can't fully participate in Microsoft 365 group services - even if the dynamic rule matches.
Azure AD will still technically add the user to the group object (visible in Azure AD), but the user will not have service-level access (Teams, SharePoint, Outlook).
In most Microsoft exam scenarios (as per the AZ-104 official study guide), such a case is treated as:
User1: added to Microsoft 365 group (licensed).
User2: skipped or treated as non-member because license missing.
Final Effective Memberships (Per Microsoft AZ-104 Study Context):
User
Group Membership
Reason
User1
Group1 only
Meets city rule; licensed user for M365; department rule may not apply because of "Human resources" not
"HR".
User2

Group1 and Group2 only
Matches both rules but has no license, still counted logically under dynamic groups in AAD object view.
Final Verified Answer (Microsoft Azure Documentation-Based):
User1: # Group1 only
User2: # Group1 and Group2 only
Microsoft Learn Extract (Supporting Evidence):
"Dynamic group membership rules automatically manage users in Azure AD based on attributes.
String comparisons are case-insensitive and must match exactly.
Assigned groups require manual membership configuration."
(Source: Microsoft Learn - Manage dynamic groups in Azure Active Directory)

NEW QUESTION: 70
□□ □□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

Name	Kind	Region
storage1	StorageV2	Central US
storage2	BlobStorage	West US
storage3	BlockBlobStorage	West US
storage4	FileStorage	East US

App1□□□ □ □□ West US Azure □□□ □□□□□. App1□ □□□□ □□□. □□□□ □□□ □□□□□ □□□. □□ □□□□ □□ □□□□ □□□ □□□□ □□□?
A. □□□2
B. □□□3
C. □□□□1
D. □□4
Answer: A (LEAVE A REPLY)

NEW QUESTION: 71
Azure Storage □□□ □□□□.
□□□□ □□□ 10□□ Blob □□□□□ □□□ □□□□□.
□□□□ □ □□□ □□, □□ □□ □□□□ □□□□□ □□ □□ □□ □□□□ □□□.
□□□□□ □□□ □□ □□□ □□ □□□?
A. □□ TLS □□□ □□□□□.
B. □□□ □□□ □□□□□.
C. □□ □□□ □□(SAS)□ □□□□□.
D. □□□ □□ □□□□□.
Answer: B (LEAVE A REPLY)
Azure Storage encrypts all data at rest using Storage Service Encryption (SSE) by default. However, if you need to use a different encryption key for specific containers or blobs within the same storage account, you must create an encryption scope.
An encryption scope defines a boundary within the storage account where data encryption is handled by a unique customer-managed key (CMK) or Microsoft-managed key. You can then associate this encryption scope with a specific container or even individual blobs, allowing flexible key management across data sets.
The Microsoft Azure Storage documentation explains that encryption scopes allow you to:
Use different encryption keys for different containers or blobs.
Rotate or manage keys independently for compliance or separation-of-duty requirements.

Support encryption at container creation by assigning a specific scope.
Other options are incorrect:
Modifying the TLS version affects network security, not encryption keys.
Generating a SAS defines access tokens, not encryption behavior.
Rotating access keys re-generates account keys for security, but it does not create a new encryption key for specific containers.
Hence, before creating the container, you must first create an encryption scope and then assign it to that container to ensure it uses a different encryption key.

NEW QUESTION: 72
□□ □□□ □□□ □□□□ Subscription1□□□ Azure □□□ □□□□.
□□: RG1
□□: □□ □□
□□: "tag1": "value1"
□□ □□□ □□□□ Policy1□□□ Azure □□□ Subscription1□ □□□□□.
□□ □□: □□
□□ □□: □□ □□ □ □□□
□□□: Policy1
□□□□:
- □□ □□: Tag2
- □□ □: Value2
Policy1□ □□□ □ □□ □□□ □□ □□□□ □□□ □□□□.
□□: storage1
□□: □□ □□
□□□ □□: RG1
□□: "tag3": "value3"
□ □□□□ □□ □□□ □□□□□□ □□□□ □□□.
□□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

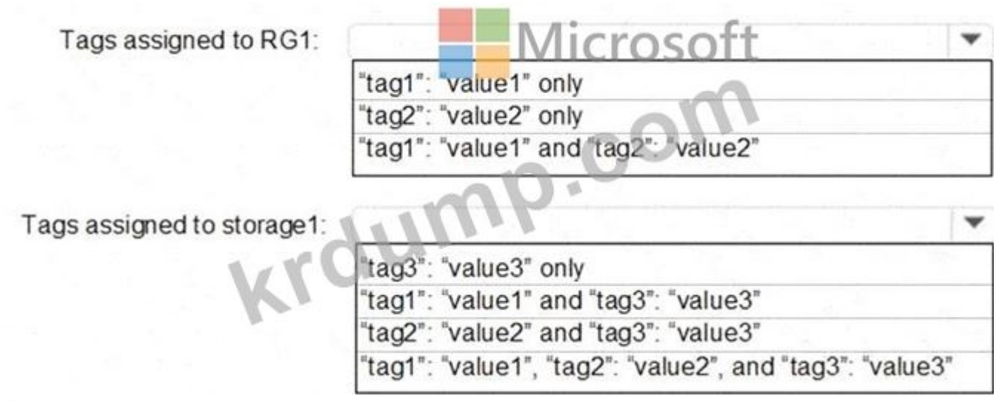

Answer:

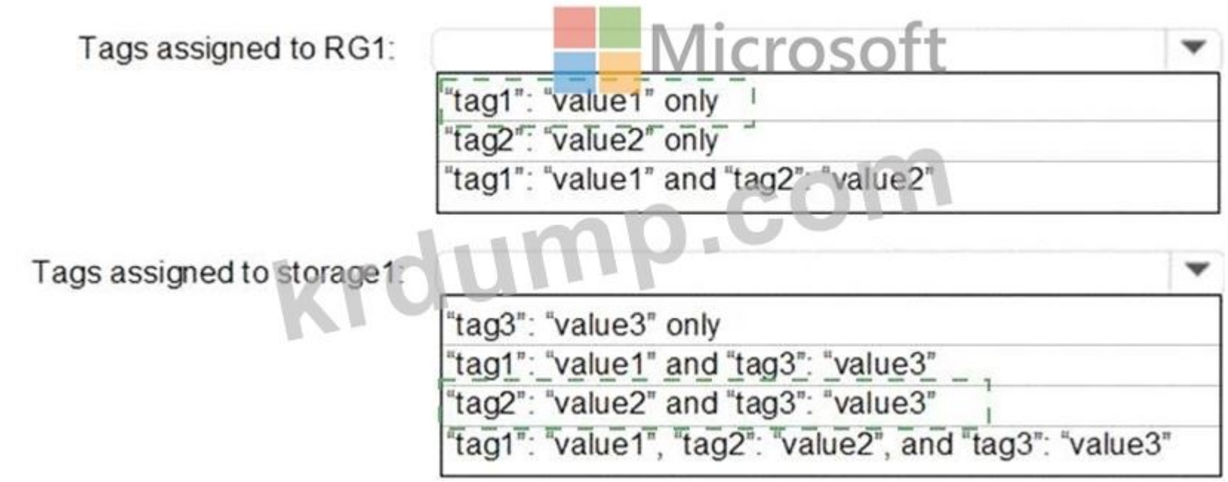

Explanation:

Box 1: "tag1": "value1" only

Box 2: "tag2": "value2" and "tag3": "value3"

Tags applied to the resource group are not inherited by the resources in that resource group.

References:

https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-using-tags

NEW QUESTION: 73

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type	Description
App1	App Service	Virtual network integration enabled for VNET1
ASP1	App Service plan	Standard SKU
VNET1	Virtual network	None
Firewall1	Azure Firewall	Connected to VNET1

Firewall1□ □□□□ VNET1□□ □□□ □□□□ □□□□ □□□.
□□ □□ □□□ □□ □□□?

A. Azure Network Watcher□ □□□□.
B. □□□□□ □□ □□□□ □□□□□.
C. ASP1□ □□□□ SKU□ □□□□□□□□.
D. □□ □□□□ □□□□□.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 74

Azure □□□ □□□□. □□ □□□□ □□ □□ □□ □□ □□ □□ □□ □□□ □□□ storage1□□□ □□□ □□□ □□□□ □□□□.

Name	Blob prefix	If base blobs were last modified more than (days ago)	Then
Rule1	container1/	3 days	Move to archive storage
Rule2	*Not applicable*	5 days	Move to cool storage
Rule3	container2/	10 days	Delete the blob
Rule4	container2/	15 days	Move to archive storage

6□ 1□□ □□ □□ □□ storage1□ □ □□ blob□ □□□□□.

Name	Location	Access tier
File1	container1	Hot
File2	container2	Hot

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
On June 6, File1 will be stored in the Cool access tier.	○	○
On June 7, File2 will be stored in the Cool access tier.	○	○
On June 16, File2 will be stored in the Archive access tier.	○	○

Answer:

Explanation:

NEW QUESTION: 75

□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Region	Peers with
VNet1	West US	VNet2
VNet2	West US	VNet1, VNet3
VNet3	East US	VNet2

□□□□ □□ □□ □□□ □□ □□□ □□□□ □□□□.

Name	Connected to
VM1	VNet1
VM2	VNet2
VM3	VNet3

□□ □□ □□□□ □□ IP □□□ □□□□.

Bastion1□□□ Azure Bastion □□□□ VNet1□ □□□□□.

Bastion1□ □□ □□ □□ □□□ □□□ □ □□□?

A. VM1□ □□

B. VM1 □ VM2□ □□

C. VM1 □ VM3□ □□

D. VM1,VM2, VM3

Answer: A (LEAVE A REPLY)

Azure Bastion provides secure and seamless RDP/SSH access to virtual machines directly through the Azure portal, without requiring a public IP address on the target VMs. Bastion is deployed per virtual network (VNet) and enables connectivity only to virtual machines within that same VNet.

In the scenario:

* Bastion1 is deployed in VNet1.

* VM1 is connected to VNet1.
* VM2 is connected to VNet2.
* VM3 is connected to VNet3.
Even though VNet1, VNet2, and VNet3 are peered, Azure Bastion access does not traverse VNet peering connections. Microsoft explicitly states:
"Azure Bastion provides RDP/SSH access to virtual machines only within the virtual network in which it is deployed. Bastion does not support access to virtual machines in peered virtual networks." Therefore, Bastion1 can only provide session connectivity to VM1, which is directly connected to VNet1.
VM2 (VNet2) and VM3 (VNet3) cannot be reached through Bastion1 even though peering exists, since Bastion connectivity is restricted to its host VNet.

NEW QUESTION: 76
contoso.com□□□ Azure Directory(Azure AD) □□□□ □□□ Azure □□□ □□□□.
□□□□ □+□□□□ Active Directory □□□□ □□□□□□. □□ □□□□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	Role
SecAdmin1	Security administrator
BillAdmin1	Billing administrator
User1	Reports reader

□□ □□□□ □□ □□ □□□ □□□□ □□□(SSPR)□ □□□□□ SSPR□ □□ □□ □□□ □□□□□ □□□□□.
□□□□ □□□ □□ □: 2
□□□□□ □□□□ □□: □□□, □□ □□
□□□ □□□ □□ □: 3
□□□□ □□□ □□ □: 3
□□ □□ □□□ □□□□□:
□□□ □□ □□□□ □□□ □□□□□?
□ □□□ □□ □□□□□?
□□□ □ □□ □□□□□ □□□ □□□□□□?
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

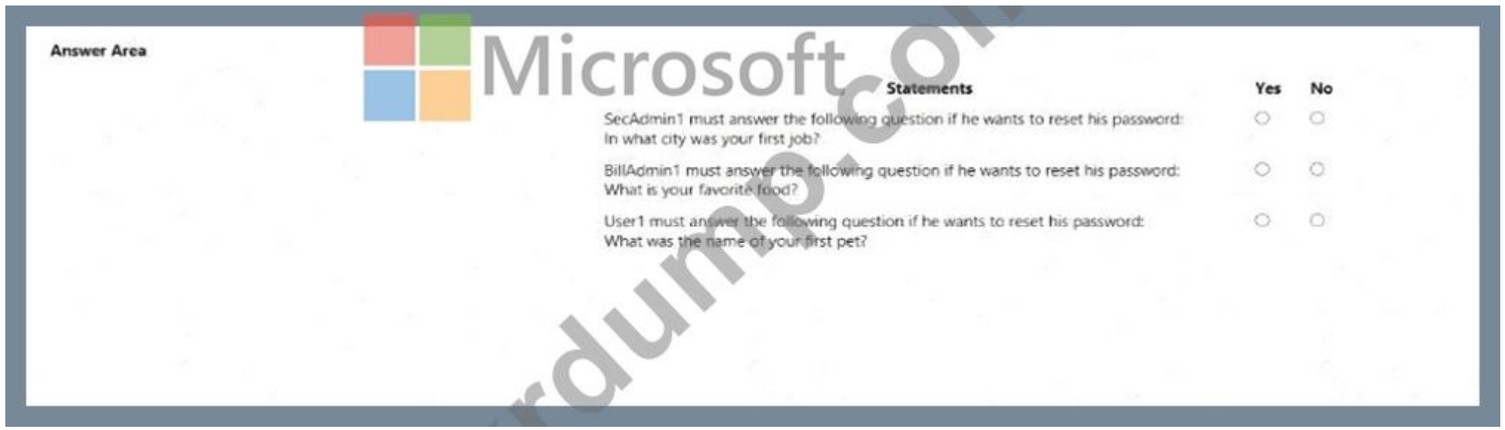

Answer:

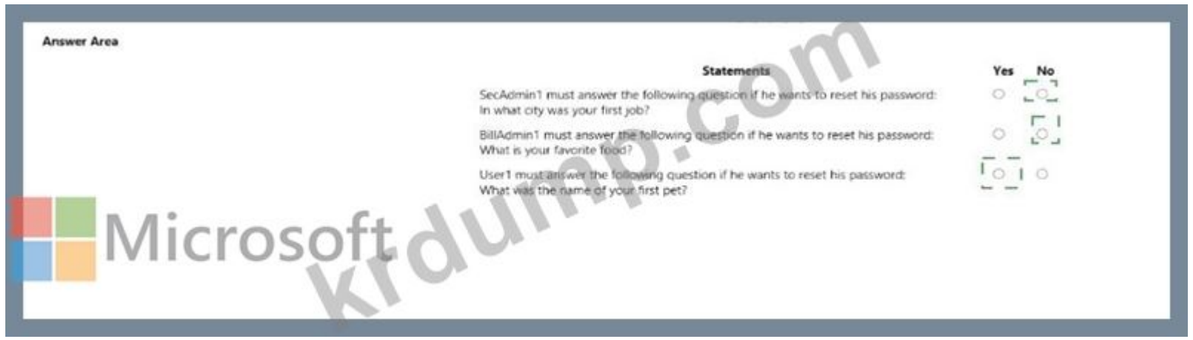

Explanation:

No, No, Yes

In Microsoft Azure Active Directory (Azure AD), the Self-Service Password Reset (SSPR) feature allows users to securely reset their passwords using preconfigured authentication methods. The configuration in this scenario specifies:

* Number of methods required to reset: 2
* Available methods: Mobile phone and Security questions
* Number of questions required to register: 3
* Number of questions required to reset: 3

However, the key detail lies in the user types and roles involved:

* SecAdmin1 (Security Administrator) and BillAdmin1 (Billing Administrator) are Azure AD administrators.
* Azure AD documentation explicitly states that administrative accounts (users with Azure AD admin roles such as Global Administrator, Security Administrator, Billing Administrator, or other privileged roles) cannot use security questions as an authentication method for SSPR.
* Admins are required to use stronger methods, such as phone, email, or app-based authentication (MFA-enabled).
* Therefore, both SecAdmin1 and BillAdmin1 will not use security questions for password reset.
* User1 (Reports reader) is a standard user, not an administrator.
* Standard users can use security questions as one of their authentication methods.
* Since security questions are enabled and three questions are required to reset the password, User1 must answer all three - including "What was the name of your first pet?" Thus, according to Azure AD SSPR policy as covered in the Microsoft Learn AZ-104 study guide and official documentation under "Self-service password reset for administrators" and "Authentication methods in Azure Active Directory", only non-administrative users can use security questions for SSPR.

NEW QUESTION: 77

□□ □□□ □□□ □□□ WEBPROD-AS-USE2□□ Azure □□□ □□□ □□□ Azure □□□ □□□□.

WEBPROD-AS-USE2□ 14□□ □□ □□□ □□□□□.

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer:

Explanation:

Box 1: 2
There are 10 update domains. The 14 VMs are shared across the 10 update domains so four update domains will have two VMs and six update domains will have one VM. Only one update domain is rebooted at a time.
Therefore, a maximum of two VMs will be offline.
Box 2: 7
There are 2 fault domains. The 14 VMs are shared across the 2 fault domains, so 7 VMs in each fault domain.
A rack failure will affect one fault domain so 7 VMs will be offline.

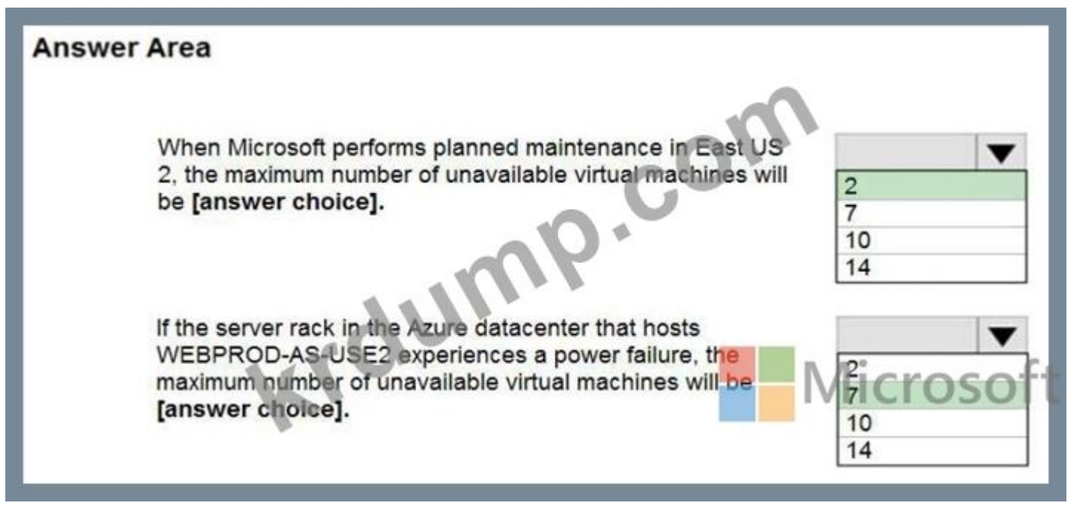

Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/manage-availability

NEW QUESTION: 78
plan1□□□ App Service □□□ webapp1□□□ Azure □□□ □□□□. □□□ □□□□ □□□ □□□ □□□ □□□ □ □□□ □□ □□ □□□□□. plan1□ □□ □□□□ □□□ □□□□ □□□.
□□ □□ □□□ □□ □□□?
A. webapp1□□ □□□□□□ □□□ □□□□□.
B. webapp1□□ □□□ □□ □□□□ □□□□□.
C. plan1□□ App Service □□□ □□□□□.
D. plan1□□ App Service □□□ □□□□□.
Answer: C (LEAVE A REPLY)
The app must be running in the Standard, Premium, or Isolated tier in order for you to enable multiple deployment slots. If the app isn't already in the Standard, Premium, or Isolated tier, you receive a message that indicates the supported tiers for enabling staged publishing. At this point, you have the option to select Upgrade and go to the Scale tab of your app before continuing.
Scale up: Get more CPU, memory, disk space, and extra features like dedicated virtual machines (VMs), custom domains and certificates, staging slots, autoscaling, and more.
Scale out: Increase the number of VM instances that run your app. You can scale out to as many as 30 instances Reference:
https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots
https://docs.microsoft.com/en-us/azure/app-service/manage-scale-up

NEW QUESTION: 79
Contoso□ □□□□ □□ □□□ □□□□ □□□.
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.

□□: □□ □□□ 1□□□□.

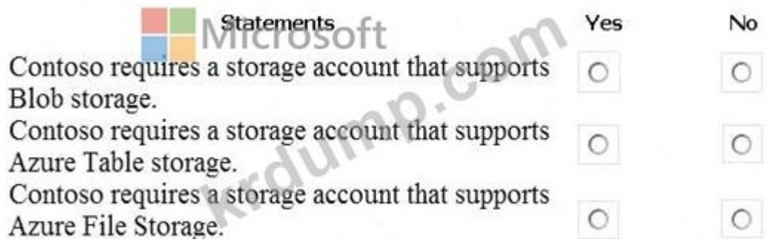

Statements	Yes	No
Contoso requires a storage account that supports Blob storage.	○	○
Contoso requires a storage account that supports Azure Table storage.	○	○
Contoso requires a storage account that supports Azure File Storage.	○	○

Answer:

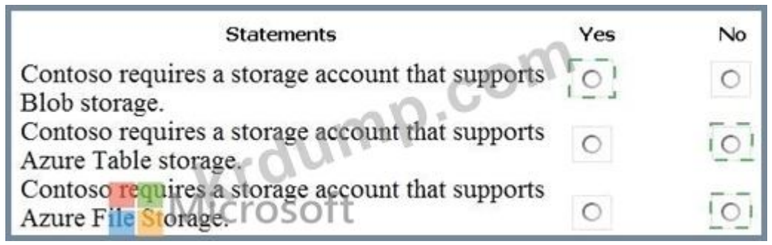

Statements	Yes	No
Contoso requires a storage account that supports Blob storage.	○	○
Contoso requires a storage account that supports Azure Table storage.	○	○
Contoso requires a storage account that supports Azure File Storage.	○	○

Explanation:

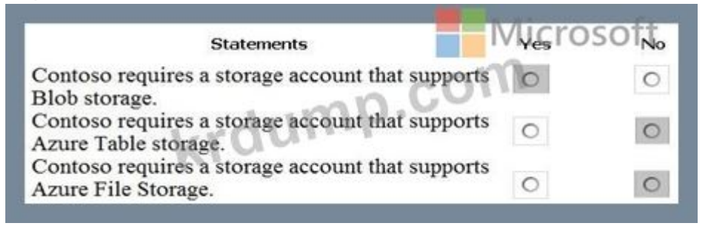

Statements	Yes	No
Contoso requires a storage account that supports Blob storage.	○	○
Contoso requires a storage account that supports Azure Table storage.	○	○
Contoso requires a storage account that supports Azure File Storage.	○	○

Statement 1: Yes

Contoso is moving the existing product blueprint files to Azure Blob storage which will ensure that the blueprint files are stored in the archive storage tier.

Use unmanaged standard storage for the hard disks of the virtual machines. We use Page Blobs for these.

Statement 2: No

Azure Table storage stores large amounts of structured data. The service is a NoSQL datastore which accepts authenticated calls from inside and outside the Azure cloud. Azure tables are ideal for storing structured, non-relational data. Common uses of Table storage include:

1. Storing TBs of structured data capable of serving web scale applications
2. Storing datasets that don't require complex joins, foreign keys, or stored procedures and can be denormalized for fast access
3. Quickly querying data using a clustered index
4. Accessing data using the OData protocol and LINQ queries with WCF Data Service .NET Libraries Statement 3: No File Storage can be used if your business use case needs to deal mostly with standard File extensions like *.

docx, *.png and *.bak then you should probably go with this storage option.

Reference:

https://docs.microsoft.com/en-us/azure/machine-learning/team-data-science-process/move-data-to-azure-blob- using-azure-storage-explorer

https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-overview

https://www.serverless360.com/blog/azure-blob-storage-vs-file-storage

NEW QUESTION: 80

VM1□□□ Azure □□ □□□ Vault1□□□ Recovery Services □□ □□ □□□ □□□□.
□□□ □□ □□ Policy1□ □□□□□. ('□□' □□ □□□□□.)

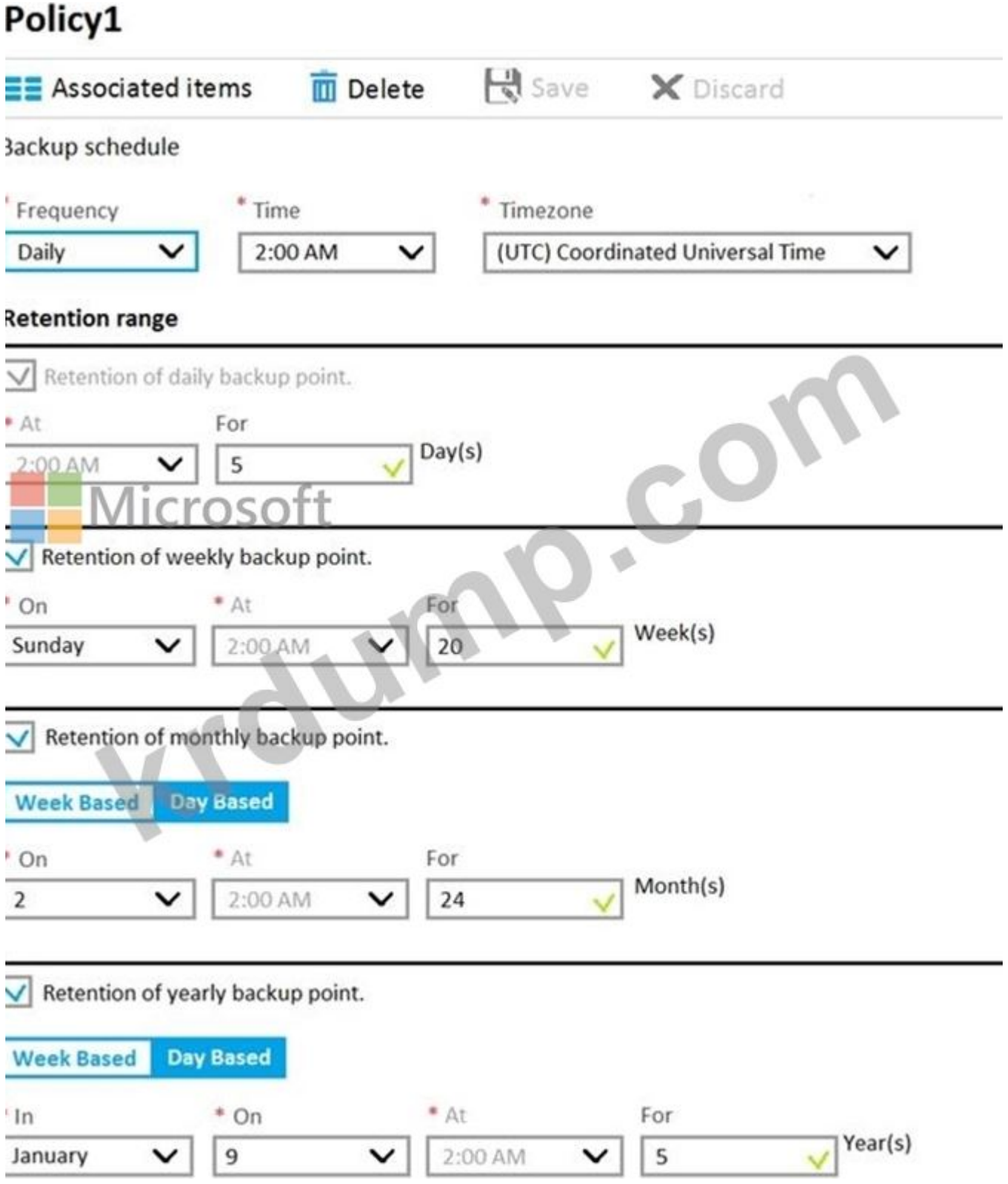

1□ 1□ □□□□ Policy1□ □□□□ VM1□ □□□ □□□□□.
VM1□ □□ □□□ □□ □□□ □□ □□□□ □□□.
1□ 8□□ 1□ 15□□ □□ □□□ □□ □□□ □ □□□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

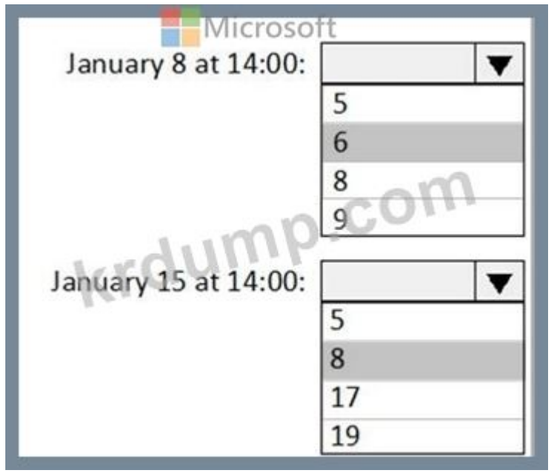

NEW QUESTION: 81
□□□ □□□□□ □ □□□□ □□ □□□ □□ □□ □□□ □□□□. □□□ □□□□ □ □□ □□□ □□ □□□ □□□□□.
Recovery Services □□ □□ □□□ □□□□ □□□.
□□ □□□ □□ □□□?
A. □□ □□□ □□ □□ □□□□ □ □□ □□□ □□□ □□□□□.
B. □□ □□□ □□ □□ □□□□ □□ □□□□ □□□□□.
C. □ □□ □□□ □□ □□ □□□ □□□□□.
D. □ □□ □□□ □□□ □□□□□.
Answer: A (LEAVE A REPLY)
You can't delete a Recovery Services vault if it is registered to a server and holds backup data. If you try to delete a vault, but can't, the vault is still configured to receive backup data.
Remove vault dependencies and delete vault
In the vault dashboard menu, scroll down to the Protected Items section, and click Backup Items. In this menu, you can stop and delete Azure File Servers, SQL Servers in Azure VM, and Azure virtual machines.

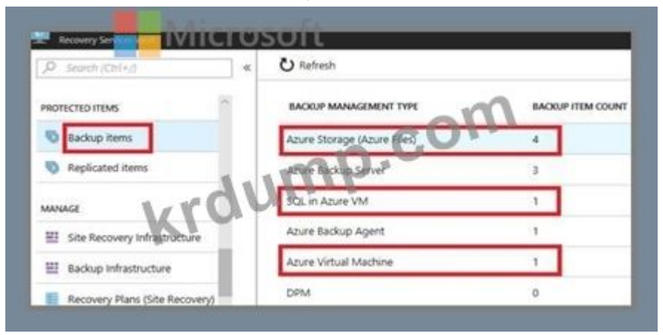

References: https://docs.microsoft.com/en-us/azure/backup/backup-azure-delete-vault

NEW QUESTION: 82
VM1□□□ Azure □□ □□□ □□□□.
Azure Backup□ □□□□ Backup1□□□ □□□ VM1 □□□ □□□□.

Backup1□ □□ □ VM1□ □□□ □□ □□ □□□ □□□□□.
VM1□ □□□ □□□□□.
Budget.xls□□ □□□ Data□□ □□□ □□□□□.
□□ □□ □□□ □□□ □□□□□ □□□□□□.
VM1□ □□□ □□□□ □□□□□.
□□□□ □□ □□ □□□ □□□ □□□□ Backup1□□ VM1□ □□□□□.
VM1□ □□ □□ □□□ □□□□□□ □□□□ □□□.
□□ □□□ □□ □□□□ □□□?
A. VM1□ □□□ □□□□□.
B. □□□ □□□□ □□□□□.
C. □□ □□□ □□□ □□□□□ □□□□□□.
D. Budget.xls□ Data□ □□□□□.
Answer: B (LEAVE A REPLY)
When you perform a VM restore using Azure Backup with the "Replace existing" option:
* Azure Backup replaces the existing VM's operating system disk and configuration with what was captured at the time of the backup.
* Data disks added after the backup are not part of the recovery point and are therefore not restored.
Other post-backup changes:
* VM size change: Automatically reverted to the size captured in the backup configuration.
* File addition (Budget.xls): Not present, because only data up to the backup time is restored.
* Password reset: Resetting password does not persist because OS disk from backup overwrites current one.However, when restoring, you can always reset credentials later - but that is expected behavior.
The only configuration that requires manual re-creation after restore is any new data disk added after the backup snapshot.

NEW QUESTION: 83
Microsoft Entra □□□□ External User□□ □□□□ □□□□ □□□□. External User□ extemall95@gmail.com□ □□□□ □□□□ □□□□□.
□□ □□□□ contractor@gmail.com□ □□□□ □□□□ □□□□□□ □□□□ □□□.
□□ □□□□□□ □□ □ □□ □□□ □□□□ □□□? □□□□ □□ □□□□ □□□ □□□ □□□□□. □□: □□□ 1□□□□.

Answer Area

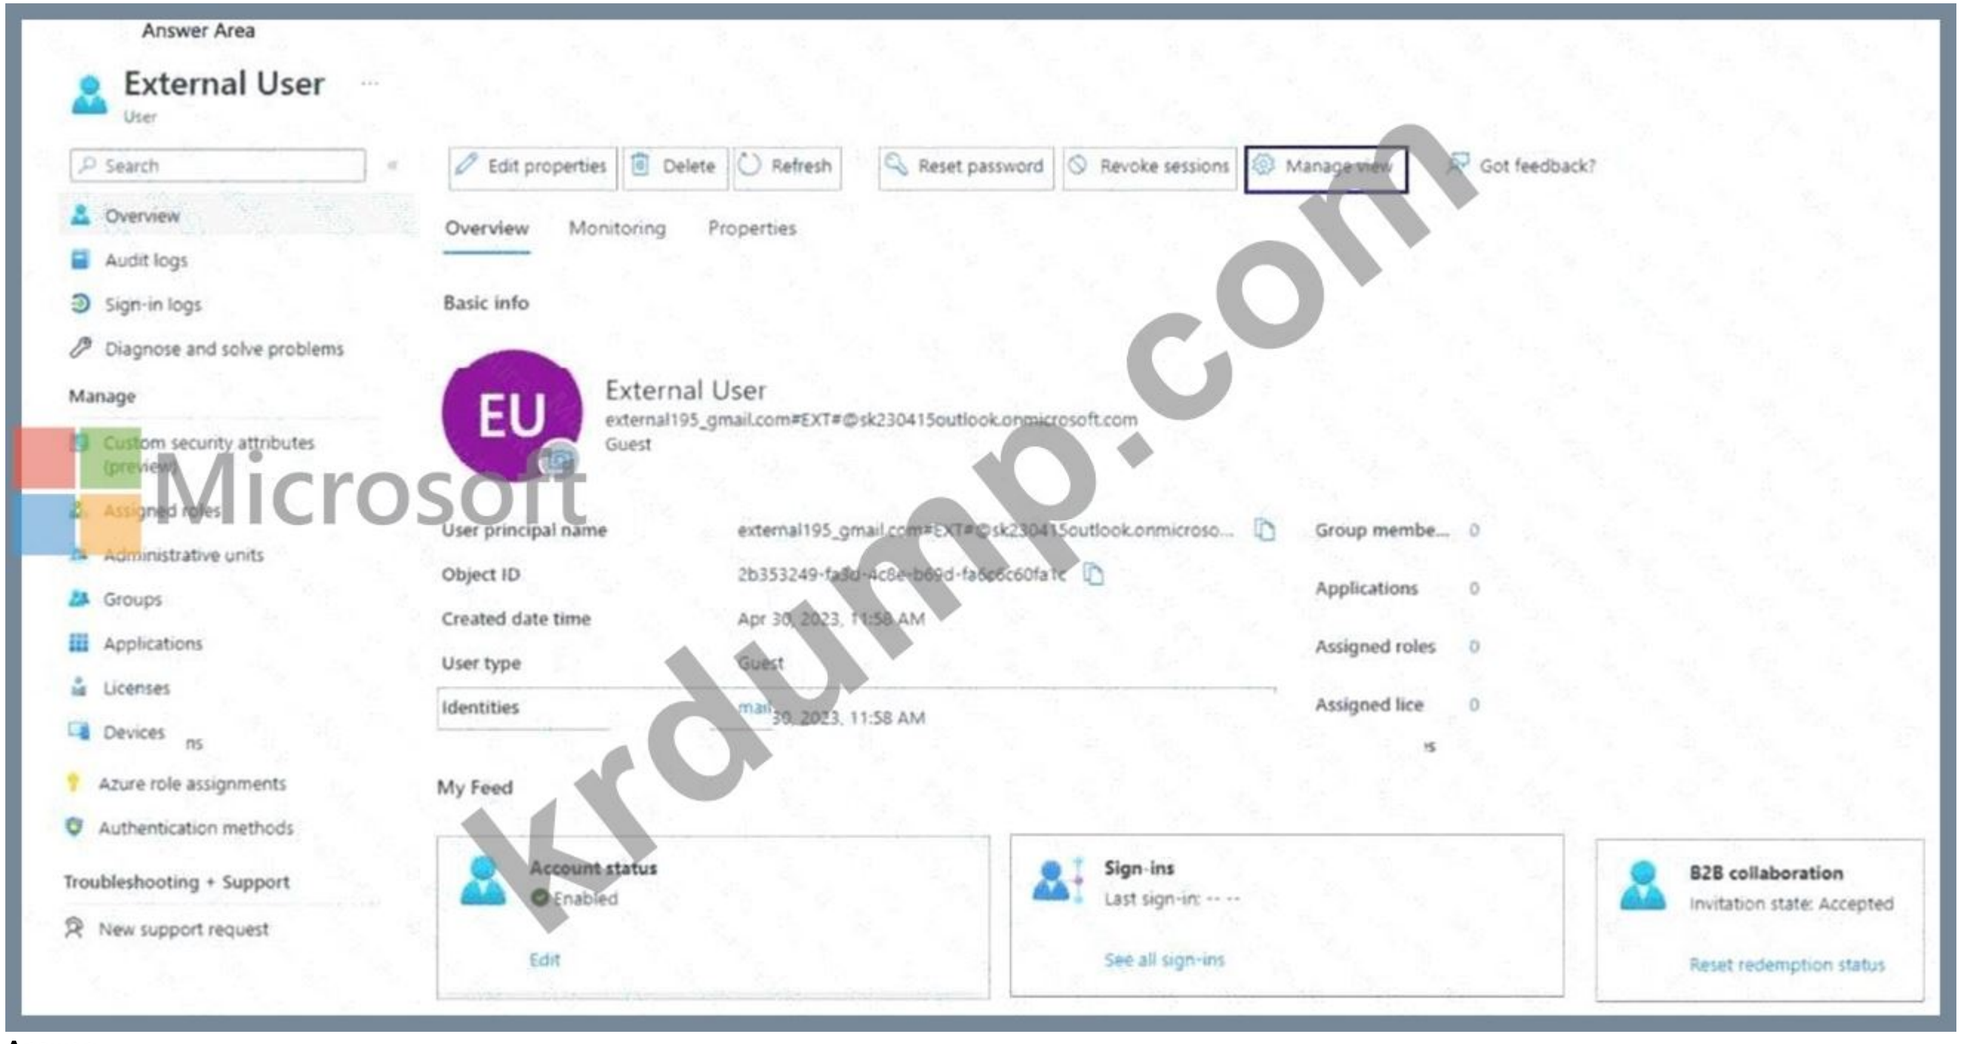

Answer:

Answer Area

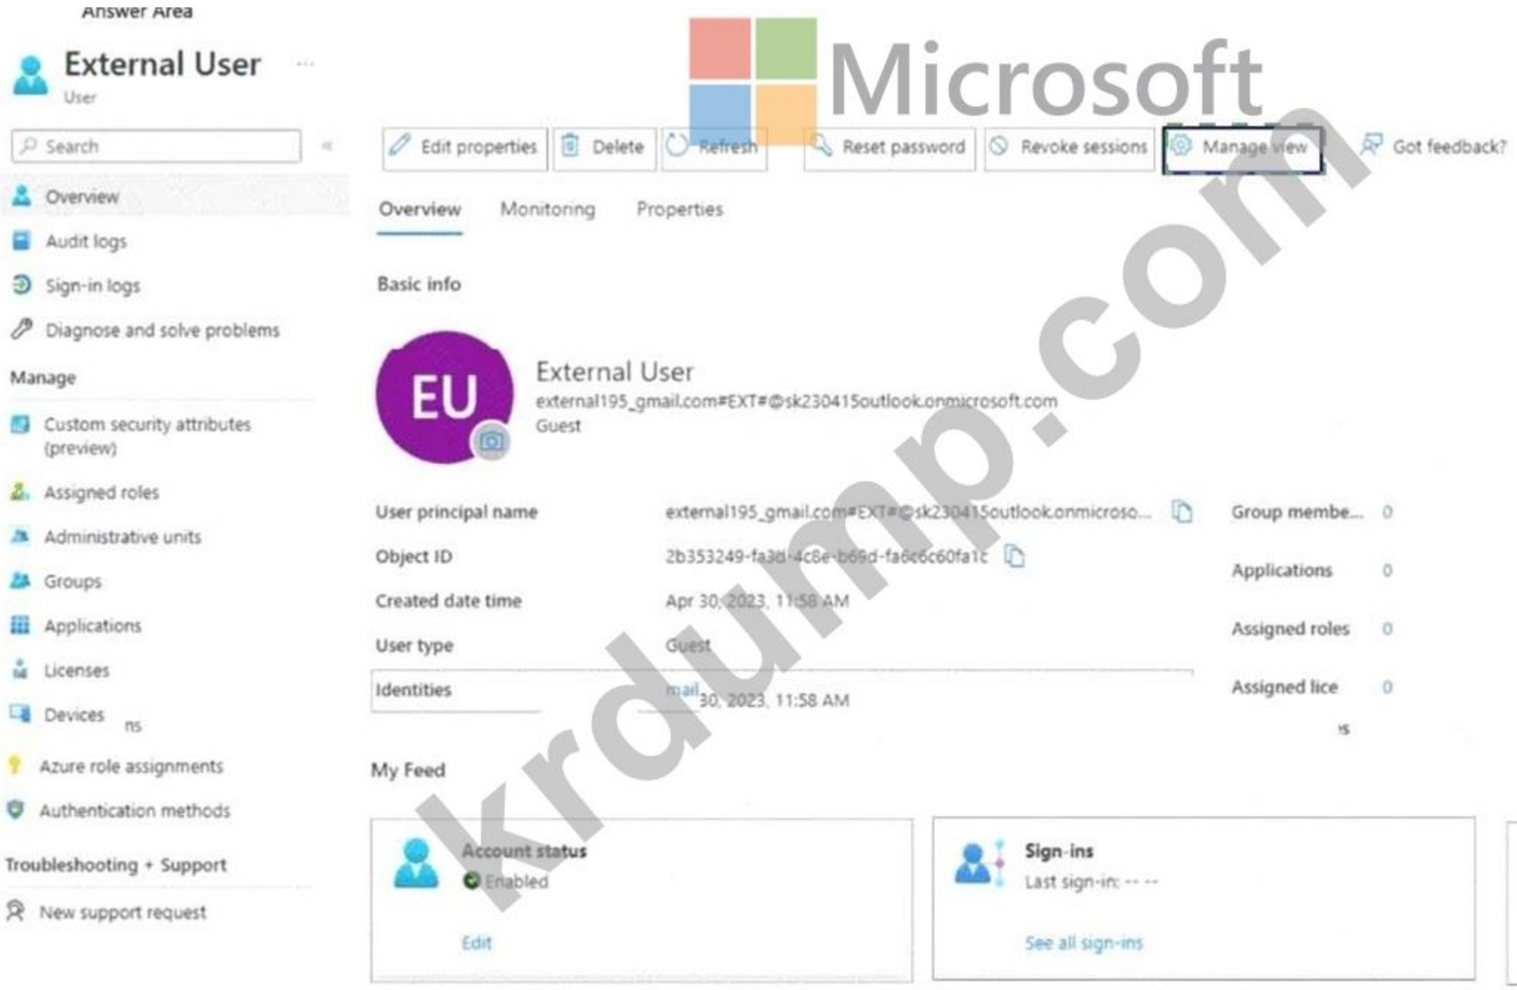

Explanation:

Answer Area

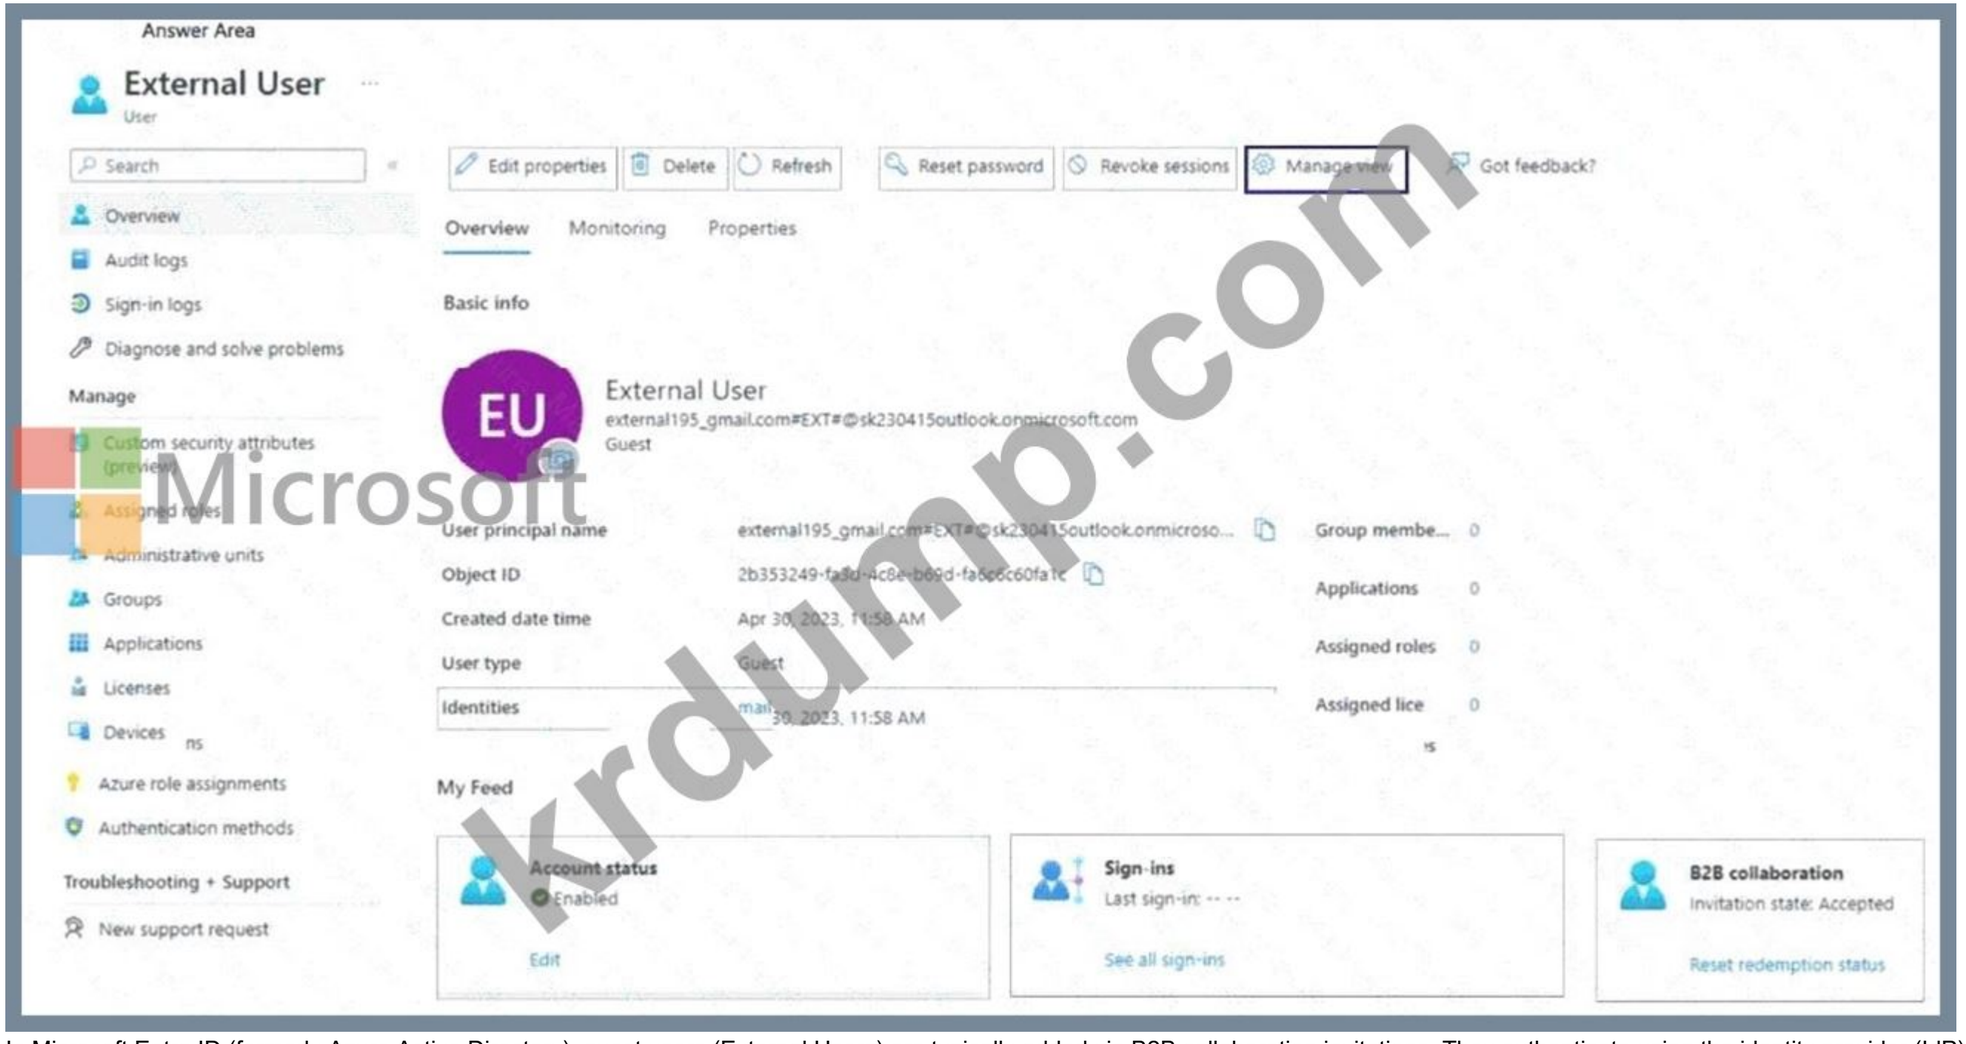

In Microsoft Entra ID (formerly Azure Active Directory), guest users (External Users) are typically added via B2B collaboration invitations. They authenticate using the identity provider (IdP) associated with the email they were invited with - such as a Microsoft account, Gmail (Google), or another organization's Azure AD tenant.

When you need to change the authentication method or identity for a guest user - for example, from external195@gmail.com to contractor@gmail.com - you must modify their identities and ensure that their account status is enabled and synchronized with the new sign-in identity.

Step-by-step from Microsoft Documentation:

From Microsoft Learn - "Manage external collaboration settings in Microsoft Entra ID" and "Manage guest accounts in Azure AD":

Identities

Under the Overview blade of the user, the Identities section lists all linked sign-in identities for that account.

You can remove the old identity (e.g., external195@gmail.com) and add the new one (contractor@gmail.

com) as an alternate identity or replace the existing one.

This determines which external account the guest uses to authenticate to your tenant.

Account status

Once the new identity is added, ensure the Account status remains Enabled so that the user can sign in successfully.
If disabled, the user cannot authenticate even with the new email address.
This setting controls whether the guest user object is active in the directory.
Other fields such as User principal name, Object ID, or Reset password cannot directly modify the external identity since guest authentication relies on the federated identity provider of the external email domain.
Therefore, the administrator must use Identities to register or replace the authentication address, and Account status to ensure the user can log in using the new account.

NEW QUESTION: 84
DCR1□ □□ □□□ □□ □□□ □□□□ □□□. □□ □□□ □□□ □□□□ □□□?
A. WQL
B. T-SQL
C. XPath
D. KQL
Answer: (SHOW ANSWER)
The planned change specifies that you must configure a Data Collection Rule (DCR) to collect only system events with Event ID 4648 from VM2 and VM4.
A Data Collection Rule (DCR) in Azure Monitor defines how data is collected from resources, filtered, and sent to destinations like Log Analytics workspaces. To define or query this data within Azure Monitor Logs or Log Analytics, you use Kusto Query Language (KQL).
From the Microsoft Learn: Azure Monitor Logs Documentation:
"Log queries in Azure Monitor are written in Kusto Query Language (KQL), the same query language used by Azure Data Explorer."
"KQL is optimized for querying large datasets, filtering by event IDs, sources, and event types." Other options:
* WQL (WMI Query Language) - used for on-prem Windows event querying, not for Azure DCR.
* T-SQL (Transact-SQL) - used for Azure SQL Database queries, not for monitoring data.
* XPath - used in Event Viewer or XML-based event filtering, not within Azure Monitor DCR configuration.
Therefore, when you configure DCR1 to collect system events (Event ID 4648) from the specified VMs, the Kusto Query Language (KQL) is the correct and verified method to filter and process these events.
Example of a valid KQL expression for this requirement:
SecurityEvent
| where EventID == 4648
| where Computer in ("VM2", "VM4")
This aligns with the Azure Monitor and Log Analytics query methodology covered in AZ-104 official exam guide (Implement and manage monitoring).

NEW QUESTION: 85
storage 1□□□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.
□□□ □□□□□ □ □□□□□ □□□□1□ □□ □□□□ □□□□ □□□. □□□□ □□ □□□ □□□□□ □□□.
□□□ □□ □□ □□ □□□ □□ □□□?
A. □□ □□□□□□ □□□□□.
B. □□ □□□□ □□□ □□□ □□□□□.
C. □□□ □□□□ □□□□□
D. □□ □□(□□ 1□) □□□ □□□□□.
Answer: B (LEAVE A REPLY)
To restrict access to a storage account in Azure to only specific networks or IP ranges (such as your home office), you must configure the firewall and virtual network settings under the networking section of the storage account.
Azure storage accounts, by default, are accessible from all networks. You can change this by modifying the Public network access setting to limit connections.

Here's the process per Microsoft Docs:
* Go to your Storage account # Networking # Firewalls and virtual networks.
* Under Public network access, select:
* "Enabled from selected virtual networks and IP addresses"
* Add:
* The IP address range of your home office.
* Any virtual networks (VNets) that should have access.
This allows you to control inbound traffic over Microsoft's backbone network without deploying additional infrastructure like Private Endpoints, which would increase administrative overhead.
Private endpoints (Option A) would indeed restrict access to a VNet via a private IP but require additional DNS and network configuration - hence not the minimal-effort option.
Access Control (IAM) (Option D) manages identity-based permissions (RBAC), not network-level access.
Therefore, modifying Public network access settings achieves the requirement with minimal effort.
Final Verified Answer: B. Modify the Public network access settings

NEW QUESTION: 86
□□ □□ □□□ Azure □□ □□□ □□□ Azure □□□ □□□□.

Name	Operating system	Subnet	Virtual network
VM1	Windows Server 2019	Subnet1	VNET1
VM2	Windows Server 2019	Subnet2	VNET1
VM3	Red Hat Enterprise Linux 7.7	Subnet3	VNET1

□□ □□ □□□ □□□ □□□□□ □□ □□□ □□□□ □□□□□□ □□□□□.

Name	DNS server
VM1	*None*
VM2	192.168.10.15
VM3	192.168.10.15

VNET1□ □□□□ □□ □□□ □□□ DNS □□□ □□□□□.

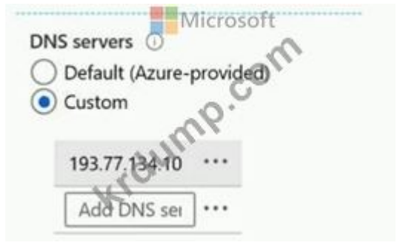

□□ □□□ IP □□□ 192.168.10.15□ DNS □□□ IP □□□ 193.77.134.10□ DNS □□□ □□□□□ □□□ □ □□□□.
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.

Answer:

Explanation:

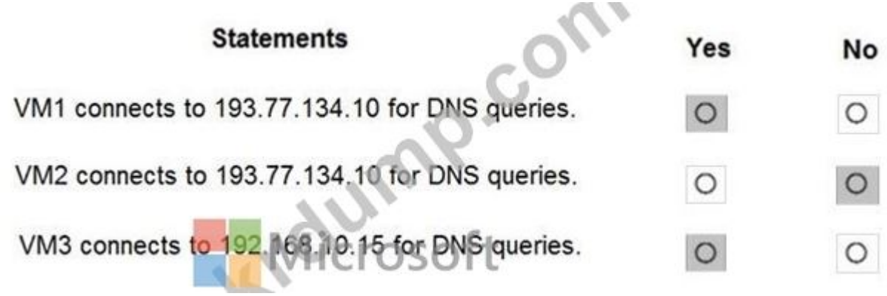

In Azure, DNS resolution order follows a specific hierarchy based on custom DNS configurations at the virtual network (VNet) and network interface (NIC) levels.

According to Microsoft Azure Administrator documentation, when you specify custom DNS servers for a virtual network, those DNS servers become the default for all virtual machines connected to that VNet - unless a specific VM's network interface has its own DNS settings configured. In that case, the NIC-level configuration overrides the VNet-level configuration.

Here's how it applies to this scenario:

VNET1 has been configured with a custom DNS server of 193.77.134.10 (visible in the VNet's DNS configuration screenshot).

VM1 does not have a custom DNS server specified on its NIC, so it inherits the VNet-level DNS server (193.77.134.10).

VM2 and VM3 both have explicit custom DNS settings pointing to 192.168.10.15 at the NIC level. Since NIC-level settings take precedence over the VNet's configuration, these VMs will use 192.168.10.15 for all DNS queries.

This hierarchy is described in Azure documentation under:

"DNS name resolution for virtual machines in Azure virtual networks," which explains that NIC-specific settings override VNet settings, and if no custom DNS is defined at either level, Azure's default internal DNS (168.63.129.16) is used.

In summary:

VM1 inherits DNS = 193.77.134.10 (Yes)

VM2 uses custom DNS = 192.168.10.15 (No for 193.77.134.10)

VM3 uses custom DNS = 192.168.10.15 (Yes)

Thus, the correct and Microsoft-verified answers are:

VM1: Yes, VM2: No, VM3: Yes

NEW QUESTION: 87

□□ □□□ □□ Policy1□□□ □□□ Recovery Services □□ □□ □□□ □□□□.

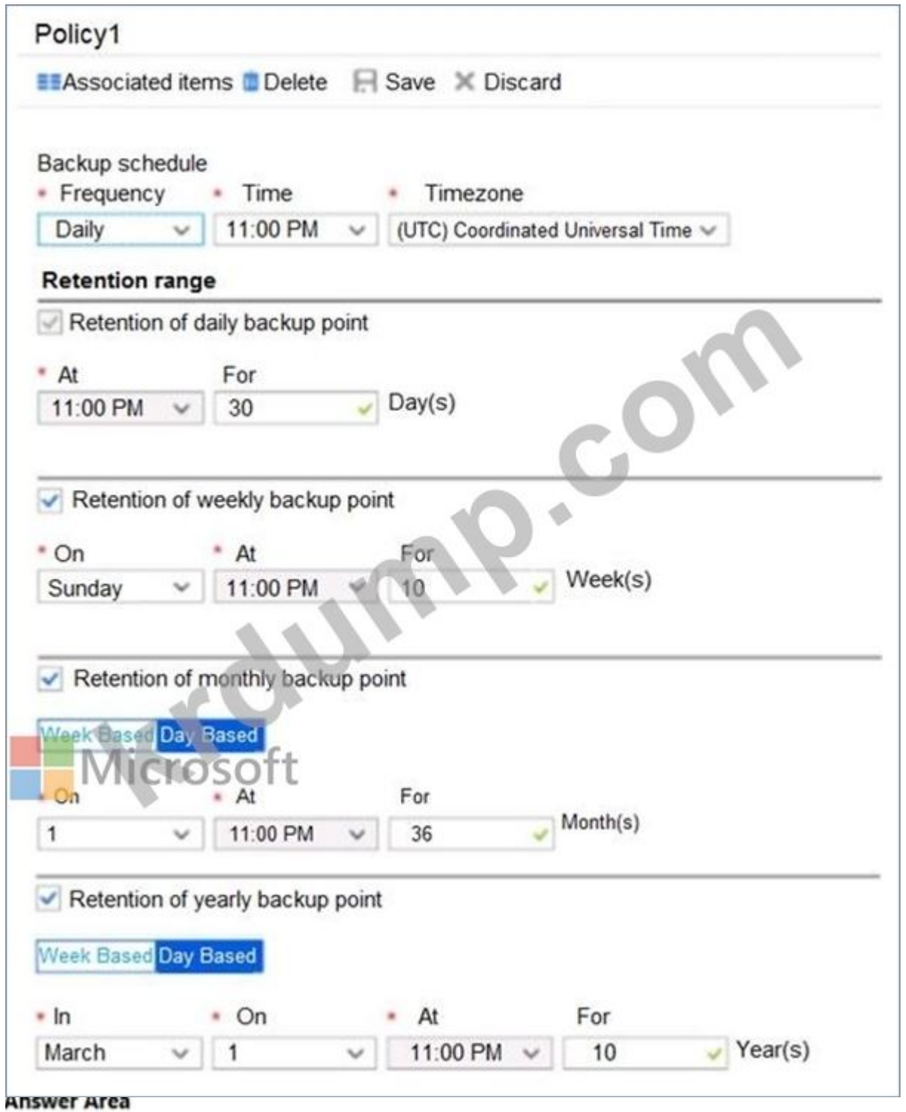

Answer Area

The backup that occurs on Sunday, March 1, will be retained for **[answer choice]**.

30 days
10 weeks
36 month
10 years

These are the selections for the statement The backup that occurs on Sunday, March 1, will be retained for [answer choice].

The backup that occurs on Sunday, November 1, will be retained for **[answer choice]**.

30 days
10 weeks
36 months
10 years

Answer:

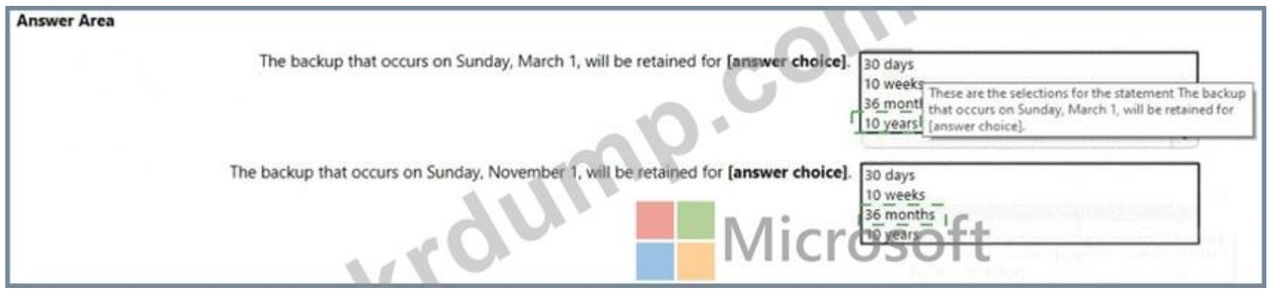

Explanation:

Box 1: 10 years

The yearly backup point occurs to 1 March and its retention period is 10 years.

Box 2: 36 months

The monthly backup point occurs on the 1

of every month and its retention period is 36 months.

Note: Azure retention policy takes the longest period of retention for each backup. In case of conflict between

2 different policies.

Reference:

https://docs.microsoft.com/en-us/microsoft-365/compliance/retention?view=o365-worldwide

NEW QUESTION: 88

storage1□□□ Azure Storage □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.

* Azure Data Lake Storage□ □□□□□.

* □□ □□□ □□ □□□□ □□ □□□ □□□□□□.

* □□□□ □□ Azure □□□ □□□□ □□□□□.

□□□□1□ □□ □□ □ □□ □□□ □□□□ □□□? □ □□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

A. □ □□□ □□

B. □ □□□ □□

C. □□□ □□□□□□

D. □□ □□ □□□(ZRS)

E. □□□ □□ □□□(GRS)

Answer: A,C,E (LEAVE A REPLY)

To meet the specified requirements for creating an Azure Storage account that supports Azure Data Lake Storage, minimizes costs, and automatically replicates to a secondary region, the following settings are appropriate:

1## Support Azure Data Lake Storage # Enable Hierarchical Namespace

Azure Data Lake Storage Gen2 requires the Hierarchical Namespace feature to be enabled in the storage account. This allows directory-based organization and fine-grained access control for big data analytics.

Answer: C. hierarchical namespace

2## Minimize costs for infrequently accessed data # Cool Access Tier

Azure Storage provides Hot, Cool, and Archive access tiers.

Hot = optimized for frequent access.

Cool = optimized for infrequent access (lower storage cost, higher access cost).

Archive = for offline, rarely accessed data.
For data that is infrequently accessed but still available online, Cool tier minimizes cost while maintaining accessibility.
Answer: A. Cool access tier
3## Automatically replicate data to a secondary region # Geo-Redundant Storage (GRS) Replication options include:
LRS (Locally Redundant Storage) - single-region redundancy.
ZRS (Zone Redundant Storage) - replication within one region.
GRS (Geo-Redundant Storage) - replicates data to a secondary Azure region automatically.
RA-GRS (Read-Access GRS) - allows read access to the secondary region.
For automatic replication to a secondary region, GRS is required.
Answer: E. geo-redundant storage (GRS)
Microsoft Documentation Extract Summary:
"Azure Data Lake Storage Gen2 is built on Azure Blob storage and requires the Hierarchical namespace setting to be enabled."
"The Cool access tier is designed for infrequently accessed data, offering reduced storage costs."
"Geo-redundant storage (GRS) replicates data to a secondary region for durability and disaster recovery." (Source: Microsoft Learn - Create and configure Azure Storage accounts)

NEW QUESTION: 89
Vault 1□□□ Recovery Services □□ □□ □□□ □□ Azure □□□ □□□□. □□ □□□□ □□ □□ □□ □□ □□ □□□ □□□□ □□□□.

Name	Operating system	Auto-shutdown
VM1	Windows Server 2016	Off
VM2	Windows Server 2022	19:00
VM3	Ubuntu Server 18.04 LTS	Off
VM4	Windows 10	19:00

□□ □ 23:00□ □□□ □□□□□ □□□ □□□□□.
Azure Backup□ □□□□ □□ □□ □□□ □□□ □ □□□?
A. VM1, VM2, VM3 □ VM4
B. VM1 □ VM3□ □□
C. VM1□
D. VM1 □ VM2□ □□
Answer: (SHOW ANSWER)

NEW QUESTION: 90
storage1□□□ Azure Storage □□□ □□□□.
storage1□ □□ □□□□ □□ □□ □□□□ □□□□□□ User1□□□ □□□□ □□□□□ □□□.
□□ □□: User1□□ Storage Account Contributor □□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: A (LEAVE A REPLY)
The Storage Account Contributor role provides the necessary permissions to manage a storage account except for access to data itself. According to Microsoft's Azure RBAC documentation, this role includes the following actions:
Microsoft.Storage/storageAccounts/regenerateKey/action
Microsoft.Storage/storageAccounts/listKeys/action

Microsoft.Storage/storageAccounts/read
Therefore, a user assigned the Storage Account Contributor role can list and regenerate access keys for the assigned storage account.
Reference from Azure documentation (Built-in roles for Azure RBAC):
"Storage Account Contributor - Manage storage accounts, including access keys. Can't manage access to data." Hence, assigning Storage Account Contributor to User1 meets the requirement.

NEW QUESTION: 91
Azure □□□ 15□ □□□□.
Group1□□□ □□ □□□ □□□ Microsoft Entra □□□□ □□□□.
□□ Azure □□□ □□□ □□□□□.
Group1□ □□ □□□ □□□ □□□ □□ □□ □□□ □□□ □ □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
* □□ □□□ □□□ □□□□□.
* □□□ □□ □□□
□□□ □□ □□□?
A. □□□ □□ □□□ □□□ Group1□ □□ □□□ □□□ □□□ □□□ □□□ □□□□□.
B. □□ □□ □□□ □□□ □□□ Group1□ □□□□□.
C. □□ □□ □□□ □□ □□□ □□□ □□□ □□□ Group1□ □□□□□.
D. □□□ □□ □□□ □□□ Group1□ □□ □□□ □□□ □□□ □□□□□.
Answer: C (LEAVE A REPLY)
The requirement is for Group1 to manage role assignments (Azure RBAC) across existing subscriptions and future subscriptions, while applying least privilege and minimizing ongoing administration. In Azure RBAC, the permission to create, update, and delete role assignments is governed by management-plane actions under Microsoft.Authorization/roleAssignments. Microsoft's Azure Administrator documentation identifies two common roles that can manage access: Owner and User Access Administrator. Of these, User Access Administrator is the least-privileged role intended specifically to manage user access without granting full resource management permissions like Owner does.
To minimize administrative effort for both current and newly purchased subscriptions, you should assign the role at the highest scope that will automatically cover all subscriptions through inheritance. The root management group is the top of the management group hierarchy; all management groups and subscriptions roll up to it. Assigning User Access Administrator to Group1 at the root management group ensures Group1 can manage role assignments across the entire hierarchy, including subscriptions added later, without repeatedly applying role assignments per subscription or per management group. This meets both requirements: least privilege (User Access Administrator instead of Owner) and minimal administrative effort (one assignment at the root scope).

NEW QUESTION: 92
WebApp1□□□ Azure App Service □□ □□□, □ □□□ Folder1□ Folder2□□ □ □□ □□□ □□□□ □□□□.
WebApp1□ □□ □□□ □□□□ □□□. □□□□ Folder2□ □□□□ □□□□□ □□ □□□.
□□ □□□ □□□□ □□, Fokier2□ □□□□□ □□□ □□□□ □□□? □□ □□□□ □□□ □□□ □□□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

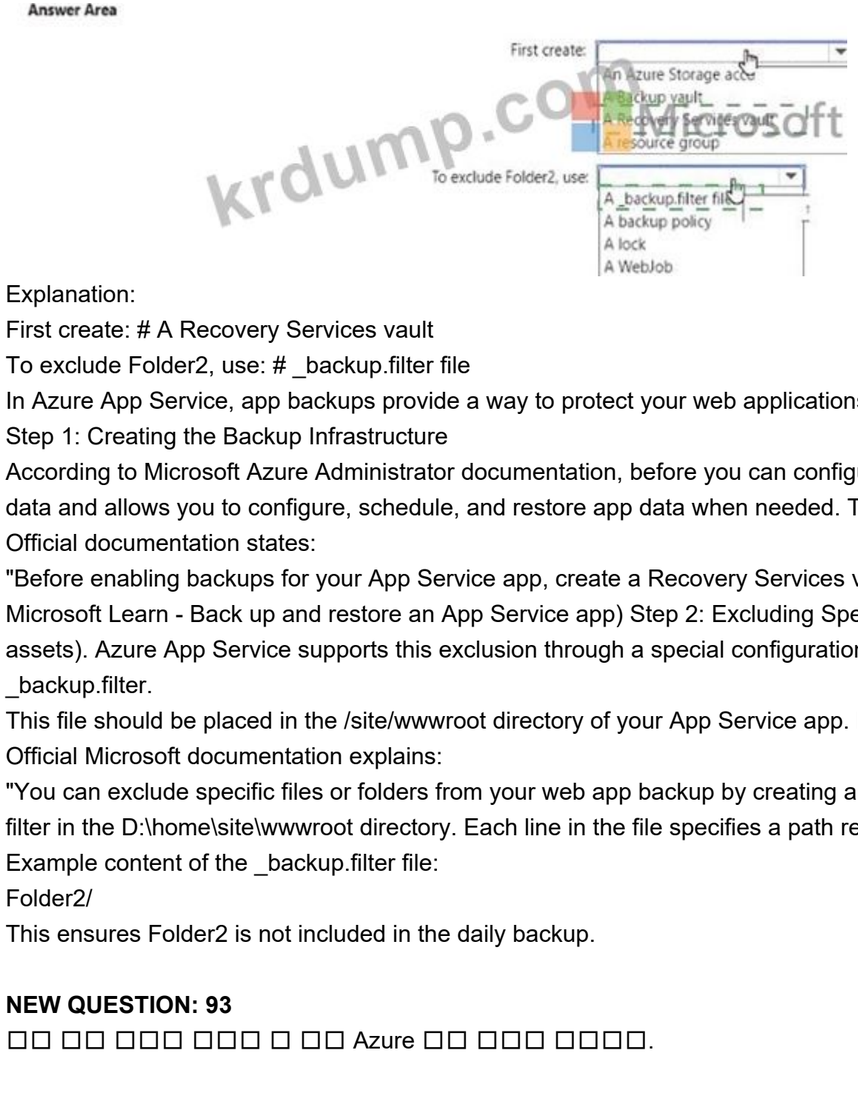

Explanation:

First create: # A Recovery Services vault

To exclude Folder2, use: # _backup.filter file

In Azure App Service, app backups provide a way to protect your web applications by storing their configuration, content, and optionally associated databases in a Recovery Services vault.

Step 1: Creating the Backup Infrastructure

According to Microsoft Azure Administrator documentation, before you can configure a web app backup, you must first create a Recovery Services vault. This vault acts as the centralized storage repository for your backup data and allows you to configure, schedule, and restore app data when needed. The vault is built on top of Azure Backup and integrates directly with App Service for automated backup scheduling.

Official documentation states:

"Before enabling backups for your App Service app, create a Recovery Services vault in the same subscription. This vault stores backup data securely and allows you to manage retention and recovery options." (Source: Microsoft Learn - Back up and restore an App Service app) Step 2: Excluding Specific Folders When performing web app backups, there might be folders you wish to exclude (such as logs, temporary files, or large static assets). Azure App Service supports this exclusion through a special configuration file named

_backup.filter.

This file should be placed in the /site/wwwroot directory of your App Service app. Inside the file, you list the relative paths (e.g., Folder2) that you want Azure to exclude during the backup process.

Official Microsoft documentation explains:

"You can exclude specific files or folders from your web app backup by creating a text file named _backup.

filter in the D:\home\site\wwwroot directory. Each line in the file specifies a path relative to the wwwroot folder that should be excluded from the backup." (Source: Microsoft Learn - Exclude files from your app backups)

Example content of the _backup.filter file:

Folder2/

This ensures Folder2 is not included in the daily backup.

NEW QUESTION: 93

□□ □□ □□□ □□□ □ □□ Azure □□ □□□ □□□□.

Name	Operating system	Private IP address	Public IP address	DNS suffix configured in the operating system	Connected to
vm1	Windows Server 2019	10.0.1.4	131.107.50.20	Contoso.com	vnet1
vm2	SUSE Linux Enterprise Server 15 (SLES) SP2	10.0.1.5	131.107.90.80	None	vnet1

□□ □□ □□□ Azure DNS □□□ □□□□.

Name	Type
Contoso.com	DNS zone
Fabrikam.com	Private DNS zone

□□ □□□ □□□□□.

fabrikam.com□ □□ vnet1□ □□ □□□□ □□□ □□□□ □□ □□□ □□□□□□.

contoso.com□ □□ vm1□ vm2□ □□□ □□□ □□□□□.

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.

□□: □□□ □□□ □□ 1□□□□ □□□□□.

Statements	Yes	No
The DNS A record for vm1 is added to contoso.com and has the IP address of 131.107.50.20.	○	○
The DNS A record for vm1 is added to fabrikam.com and has the IP address of 10.0.1.4.	○	○
The DNS A record for vm2 is added to fabrikam.com and has the IP address of 10.0.1.5.	○	○

Answer:

Statements	Yes	No
The DNS A record for vm1 is added to contoso.com and has the IP address of 131.107.50.20.	○	○
The DNS A record for vm1 is added to fabrikam.com and has the IP address of 10.0.1.4.	○	○
The DNS A record for vm2 is added to fabrikam.com and has the IP address of 10.0.1.5.	○	○

Explanation:

\

\

NO

YES

YES

NEW QUESTION: 94

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type	Resource group	Tag
RG6	Resource group	*Not applicable*	*None*
VNET1	Virtual network	RG6	Department: D1

□□ □□ □□□ □□ RG6□ □□□ □□□□□.

Section	Setting	Value
Scope	Scope	Subscription1/RG6
	Exclusions	*None*
Basics	Policy definition	Apply tag and its default value
	Assignment name	Apply tag and its default value
Parameters	Tag name	Label
	Tag value	Value1

RG6□ □□□ □□□□□: RGroup: RG6.

VNET2□□ □□ □□□□□ RG6□ □□□□□.

VNET1□ VNET2□ □□ □□□ □□□□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

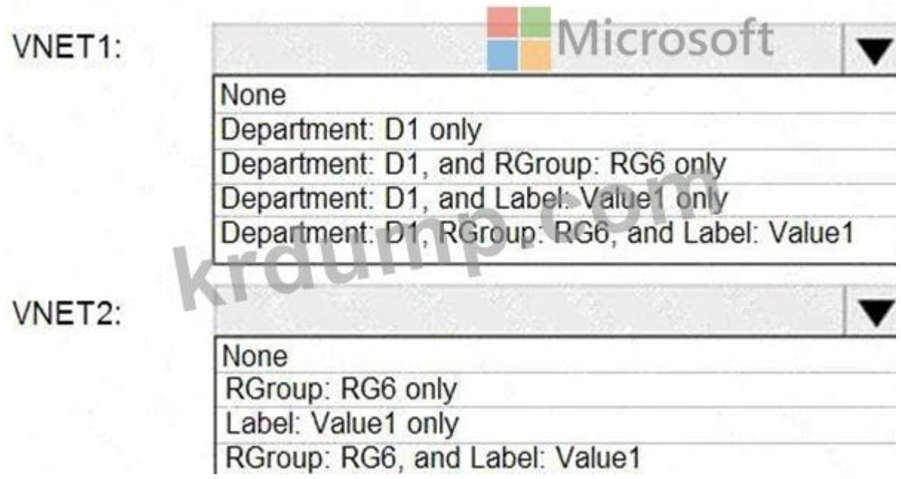

Answer:

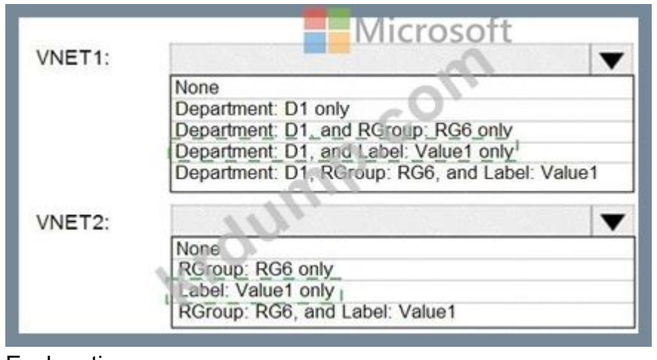

Explanation:

Answer Area

https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/tag-policies According to Microsoft Azure Resource Manager (ARM) and Azure Policy documentation used in the Microsoft Certified: Azure Administrator Associate (AZ-104) study guide, tags are name-value pairs applied to Azure resources to logically organize and manage them.
Tags can be manually applied or automatically enforced through Azure Policy. In this question, the assigned policy is "Apply tag and its default value", scoped to Resource Group RG6. The policy parameters specify that the Tag name is "Label" and the Tag value is "Value1".
Here's how Azure applies tags in this context:
* Tags applied at the resource group level are not inherited by resources within the group automatically.
They must be manually set or enforced through a policy.
* When you apply the "Apply tag and its default value" policy to a resource group, it automatically assigns the specified tag ("Label: Value1") to all existing and new resources within that scope (RG6), unless the resource already has a tag with the same name.
* The existing tag on VNET1 (Department: D1) remains unchanged because the policy does not overwrite existing tags-it only adds the new tag if it's missing.
* VNET2, which is deployed after the policy is assigned, inherits the "Label: Value1" tag automatically because it's a newly created resource within RG6.
Since RG6 itself has a tag "RGroup: RG6", that tag does not automatically apply to its resources because tag inheritance does not occur from resource groups unless enforced by a policy.
Therefore:
* VNET1 retains its original tag "Department: D1" and receives the additional "Label: Value1" tag from the applied policy.
* VNET2 only receives "Label: Value1" because that's the policy-enforced tag for the resource group RG6.
Final Verified Answers:
VNET1: Department: D1 and Label: Value1 only
VNET2: Label: Value1 only

NEW QUESTION: 95
www.contoso.com□ □□□□□□ contoso.azurewebsites.net□□□ Azure □□□ □□□□ □□□.
□□ □□ □□□ □□ □□□?
A. □□□ □□ ID□ □□□□ asuid□□ CNAME □□□□ □□□□.
B. www.contoso.com □ asuid.contoso.com□□□ □□□ A □□□□ □□□□.
C. □□□ □□ ID□ □□□□ asuid□□ TXT □□□□ □□□□.
D. www.contoso.com□□□ □□□ TXT □□□□ □□□ □□ contoso.azurewebsites.net□□□.
Answer: (SHOW ANSWER)
When you configure a custom domain (like www.contoso.com) for an Azure Web App (App Service), Azure requires that the domain be verified to ensure ownership before binding it to the app.
According to Microsoft Azure official documentation ("Map a custom domain name to your Azure web app"
- Microsoft Learn):
"Before you can add a custom domain, you must verify that you own the domain name by creating a DNS record with your domain registrar. Azure uses an asuid verification record, which can be either a CNAME or TXT record, depending on your DNS provider." The first step is to create a CNAME or TXT record in your DNS zone that links your custom domain to the Azure verification ID.
The record name is asuid.contoso.com.
The value (target) is the domain verification ID shown in the Azure portal under Custom domains # Custom hostnames # Domain ownership.
Once Azure verifies the domain ownership using that record, you can add www.contoso.com as a custom hostname in the web app configuration.
Key Point:
Step 1: Verify domain ownership using a CNAME or TXT record.

Step 2: Bind the custom hostname (www.contoso.com) to your web app.
Thus, the correct and verified answer is:
A. Create a CNAME record named asuid that contains the domain verification ID.

NEW QUESTION: 96
VM1□□□ □□ □□□ User1, User2□□ □ □□ □□□□ □□□ Azure □□□ □□□□.
User1□□ VM1□ □□□ □□□ □□□□□.
User1□ VM1□□ □□□ □ □□ □ □□ □□□ □□□□□? □ □□□ □□□ □□□□ □□□□□.
A. □□□ □□□□ □□□□□.
B. □□ □□□ □□□□□.
C. User2□□ VM1□ □□□ □□□ □□□□□.
D. VM1 □□□□ Azure □□□ □□□□ □□□□□□.
E. User2□□ VM1□ □□ Reader □□□ □□□□□.
Answer: A,D (LEAVE A REPLY)
In Microsoft Azure, Role-Based Access Control (RBAC) defines what actions users can perform on resources. The Contributor role is one of the built-in Azure roles that provides extensive permissions but excludes access management rights.
According to the Microsoft Azure Administrator documentation under "Built-in roles for Azure resources", the Contributor role has the following definition:
"Grants full access to manage all Azure resources, including the ability to create and manage resources, but does not grant permission to assign roles or manage access rights." This means that a Contributor can perform any configuration or management task on a resource (such as starting, stopping, resizing, or modifying a virtual machine) but cannot grant or modify permissions for other users.
Let's analyze each option in this context:
A). Add a data disk - # Correct
A Contributor can attach, detach, or manage data disks for the VM.
These actions involve resource management operations under the Microsoft.Compute/virtualMachines namespace, which are permitted to the Contributor role.
Microsoft Documentation Reference (Compute Resource Provider Operations):
Contributors can "create, update, and delete virtual machines, disks, and configurations." B). Configure a daily backup - # Incorrect Configuring a daily backup requires permissions within the Recovery Services vault or Backup vault resource type (Microsoft.RecoveryServices/vaults/*).
The Contributor role on the VM does not automatically grant access to vault-level resources.
Hence, User1 cannot configure or schedule a daily backup without additional permissions on the vault.
C). Assign User2 the Contributor role for VM1 - # Incorrect
Assigning or modifying RBAC roles requires the Owner or User Access Administrator role because it involves operations under the Microsoft.Authorization/roleAssignments/* permission.
The Contributor role does not include these permissions.
Therefore, User1 cannot assign roles to others.
D). Upload an image of VM1 to an Azure compute gallery - # Correct
The Contributor can capture or generalize a VM and upload its image to an Azure Compute Gallery, as this operation falls under the VM and image management permissions (Microsoft.Compute/galleries/images/*).
This action is considered a resource management operation and does not require ownership or RBAC rights beyond Contributor.
E). Assign User2 the Reader role for VM1 - # Incorrect
Similar to option C, assigning RBAC roles is a management operation restricted to users with Owner or User Access Administrator privileges.
Contributors cannot modify or assign access control permissions.

NEW QUESTION: 97
storage1□□□ Azure Storage □□□ □□, □ □□□□ □□□□ 1□ □□□□ 2□□ □ □□ □□□□□ □□□□. □ □□□□ □□□ Blob □□ □□□ □□□□□ □□□□.

□□□□□ □□□ □□□ □□ □□□□ □□□□.
□□□ □□ □□ □□ □□ □□□ □□□□.

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

Based on the lifecycle management policy you created and the information from the web search results, here are the answers to your statements:

A blob snapshot automatically moves to the Cool access tier after 15 days. = Yes A blob version in container2 automatically moves to the Archive access tier after 30 days. = No A rehydrated version automatically moves to the Archive access tier after 30 days. = No The lifecycle management policy you created has two rules: one for container1 and one for container2. The rule for container1 has an action that moves blob snapshots to the Cool access tier if they are older than 15 days. Therefore, a blob snapshot in container1 will automatically move to the Cool access tier after 15 days, regardless of the access tier of the base blob.

The rule for container2 has an action that moves blob versions to the Archive access tier if they are older than
30 days and have a prefix match of "archive/". Therefore, a blob version in container2 will only automatically move to the Archive access tier after 30 days if its name starts with "archive/". Otherwise, it will remain in its current access tier.
A rehydrated version is a blob version that was previously in the Archive access tier and was restored to an online access tier (Hot or Cool) by using the rehydrate priority option1. A rehydrated version does not automatically move to the Archive access tier after 30 days, unless there is a lifecycle management policy rule that explicitly specifies this action. In your case, neither of the rules applies to rehydrated versions, so they will stay in their online access tiers until you manually change them or delete them.

NEW QUESTION: 98
User1□ □□□□□ □□□ □□□ □ □□, User4□ RG2□ □□□□□□ □□□ □ □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
□ □□□□□ □□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

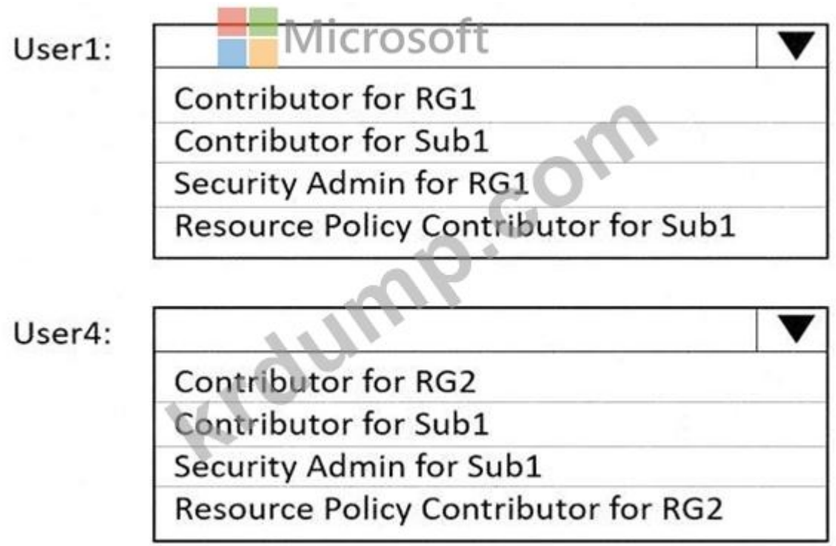

Answer:

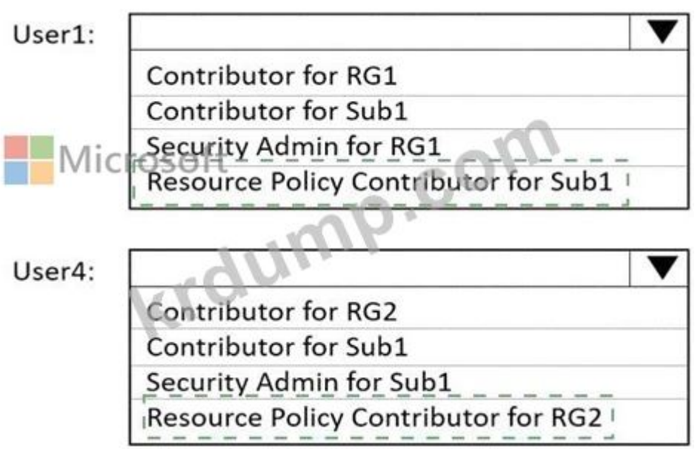

Explanation:

To meet the technical requirement that User1 can create initiative definitions and User4 can assign initiatives to RG2, we must understand how Azure Policy roles operate according to Microsoft Azure documentation and AZ-104 study guides.

Azure Policy uses role-based access control (RBAC) to determine who can create, manage, and assign policies or initiatives (policy sets). The key roles relevant to this task are:

* Resource Policy Contributor - Allows users to create, edit, and delete policy definitions and initiative definitions, but does not allow assignment of policies or initiatives. This role is required for users who will design or build the policy logic (definitions).
* Therefore, User1, who needs to create initiative definitions, must have the Resource Policy Contributor role.
* Scope: Subscription (Sub1) - because initiatives are created at subscription or management group level.
* Contributor - Provides full access to manage resources, except granting permissions. Contributors can assign existing policies or initiatives to resource groups or subscriptions if they have sufficient permissions.
* Therefore, User4, who needs to assign initiatives to RG2, must be assigned the Contributor role at the RG2 scope.

Other roles:

* Security Admin is related to Microsoft Defender for Cloud security settings, not policy creation.
* Resource Policy Contributor for RG2 would not meet the requirement because creating definitions typically requires subscription-level access.

According to Microsoft Docs ("Azure built-in roles for Policy"):

"The Resource Policy Contributor role allows users to create, edit, and delete Azure policy and initiative definitions. To assign policies or initiatives, users need the Contributor or Owner role at the appropriate scope." Hence, the configuration should be:

* User1: Resource Policy Contributor for Sub1
* User4: Contributor for RG2

NEW QUESTION: 99

□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

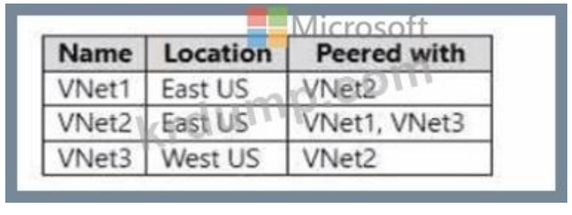

Name	Location	Peered with
VNet1	East US	VNet2
VNet2	East US	VNet1, VNet3
VNet3	West US	VNet2

□□□□ □□ □□ □□□ □□ □□□ □□□□ □□□□.

Name	Operating system	Connected to
VM1	Windows	VNet1
VM2	Linux	VNet2
VM3	Windows	VNet3

□ □□ □□□□ □□ IP □□□ □□□□□.
□□ □□□ □□ VNet1□ □□ Azure □□□ □□□□.

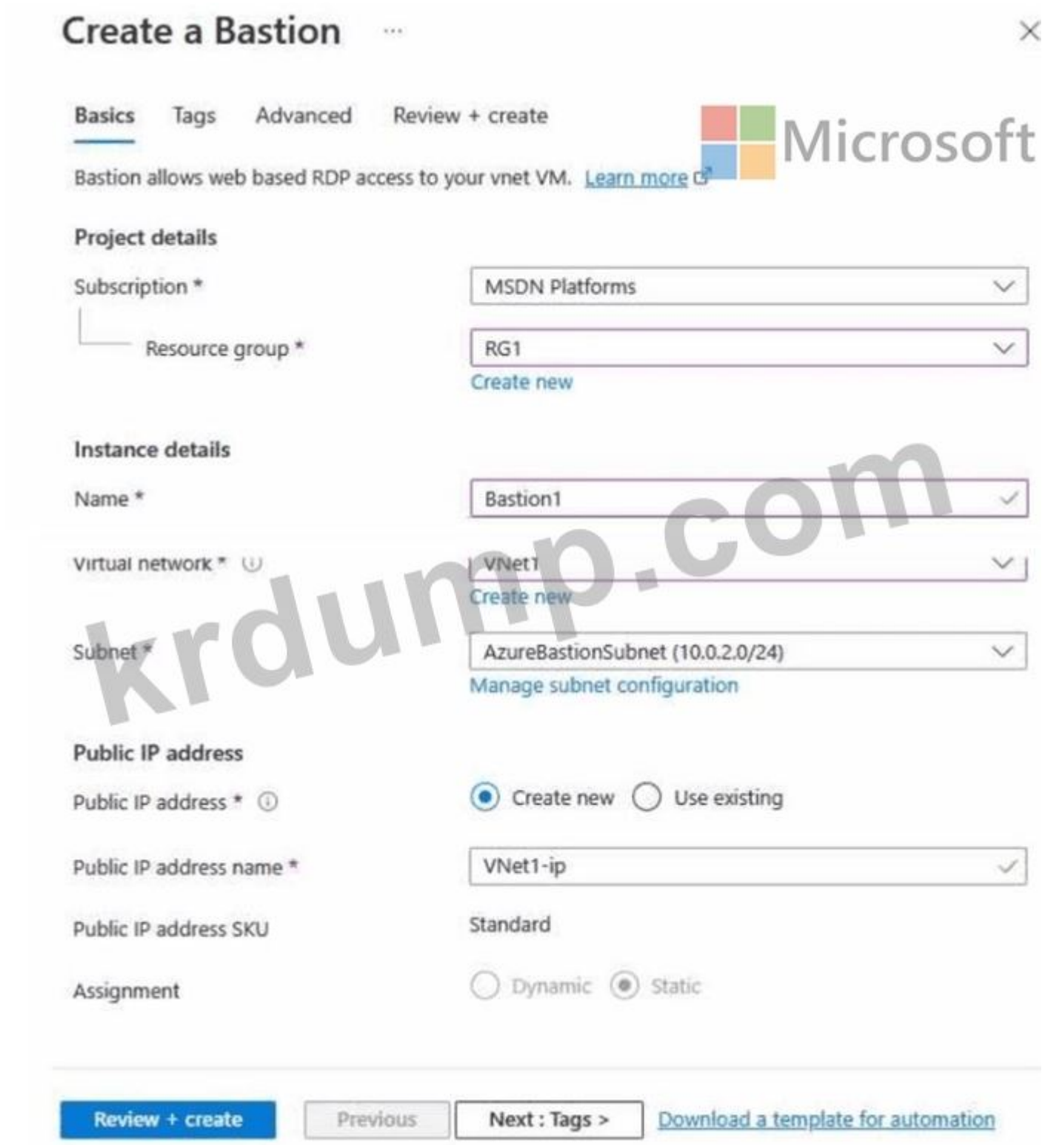

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

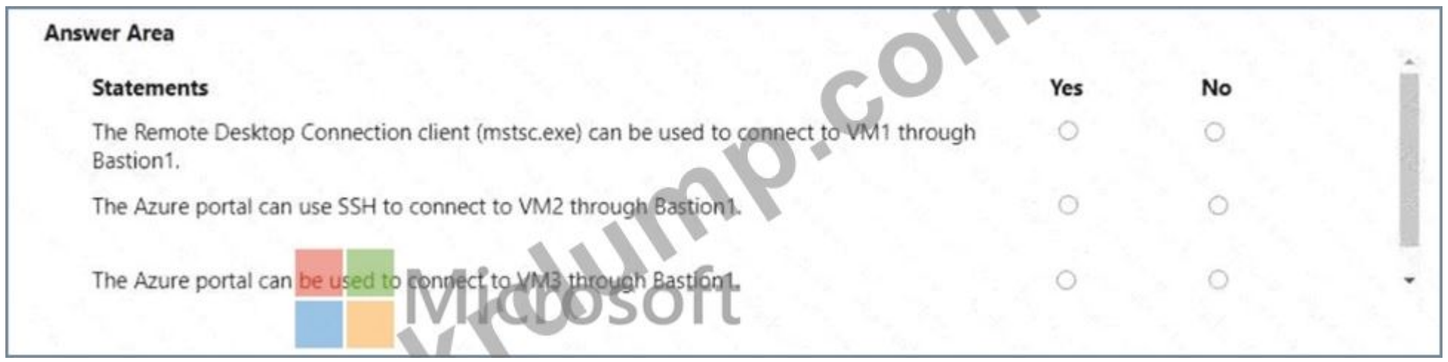

Answer:

Explanation:
No
Yes
No
Azure Bastion is a fully managed platform-as-a-service (PaaS) solution that allows secure RDP (for Windows VMs) and SSH (for Linux VMs) connections directly through the Azure portal - without exposing public IP addresses. It provides browser-based connectivity over TLS from the Azure portal and uses port 443 only.
Let's analyze the scenario step by step based on the provided configuration and the Microsoft Azure Administrator documentation (AZ-104 study guide):
Network Topology Review
VNet
Location
Peered With
VNet1
East US
VNet2
VNet2

East US
VNet1, VNet3
VNet3
West US
VNet2
VM
OS
Connected To
VM1
Windows
VNet1
VM2
Linux
VNet2
VM3
Windows
VNet3
Bastion1 is deployed in VNet1, which is peered with VNet2 (but not directly with VNet3).
Statement Analysis
1. The Remote Desktop Connection client (mstsc.exe) can be used to connect to VM1 through Bastion1 Azure Bastion does not support native clients like mstsc.exe.
Bastion connections are made only via the Azure portal using web-based RDP or SSH.
The mstsc.exe client requires a direct RDP connection, which is not provided through Azure Bastion.
Answer: No
2. The Azure portal can use SSH to connect to VM2 through Bastion1
Bastion1 is deployed in VNet1, which is peered with VNet2.
Azure Bastion supports connections to peered VNets within the same region, as long as network peering allows traffic between VNets and the VM does not require a public IP.
VM2 (Linux) is connected to VNet2, which is peered with VNet1 (same region, East US).
Therefore, the Azure portal can use SSH (port 22) through Bastion1 to connect to VM2.
Answer: Yes
3. The Azure portal can be used to connect to VM3 through Bastion1
VM3 is connected to VNet3, which is in West US.
Bastion1 (in VNet1, East US) cannot connect across regions, even though VNet2 is peered with both VNet1 and VNet3.
Bastion does not support transitive peering - meaning Bastion1 cannot connect to VMs in VNets indirectly connected (VNet3 in this case).
Answer: No
Key Azure Documentation Points (Microsoft Learn Extract):
"Azure Bastion allows you to securely connect to a VM in the same virtual network or a peered virtual network in the same region."
"Bastion does not support transitive peering or connections across regions."
"Connections are made through the Azure portal only and not via native RDP or SSH clients." (Source: Microsoft Learn - Azure Bastion Overview and Connectivity Requirements)

NEW QUESTION: 100
Vault1□□□ Recovery Services □□ □□ □□□ □□□ Azure □□□ □□□□.

Vaultl□ □□ □□ □□□ □□ □□(MAU)□ □□□□□ □□□.
□□ □□□□ □□ □□□□ □□□?
A. □□□□ ID
B. □□ □□□
C. □□ □□
D. □□□ □□ Azure □□
Answer: B (LEAVE A REPLY)
https://learn.microsoft.com/en-us/azure/backup/multi-user-authorization?tabs=azure-portal&pivots=vaults- recovery-services-vault#before-you-start Before you start Ensure the Resource Guard and the Recovery Services vault are in the same Azure region.
Ensure the Backup admin does not have Contributor permissions on the Resource Guard. You can choose to have the Resource Guard in another subscription of the same directory or in another directory to ensure maximum isolation.
Ensure that your subscriptions containing the Recovery Services vault as well as the Resource Guard (in different subscriptions or tenants) are registered to use the providers - Microsoft.RecoveryServices and Microsoft.DataProtection . For more information, see Azure

NEW QUESTION: 101
VNet1□ VNet2□□ □ □□ □□□□ □□ □□□□□ □□□ Azure □□□ □□□□. VNet1□□ □□ □□□□ □□□□ VPN □□□□□□ □□□□.
□□□□□ □□□□□□ VNet1□ VPN □□□□□□ □□□□ VPN □□□ □□□□.
VNet2□ □□ □□□ □□□□□ □□□□□ □□□□□ □□□□ □□ □□□□ □□□□ □□□.
□□□□ □□□ □□□□□ □□□.
□□ □□□ □□□ □□□□ □□□?
A. □□ IP □□□ □□ Azure □□□
B. VNet2□ □□ ExpressRoute □□
C. □□□ □□□ □ □□□ □□ □□(UDR)
D. Azure □□□□□□ □□□□□
Answer: C (LEAVE A REPLY)
Because VNet1 and VNet2 are peered, and VNet1 already has a VPN gateway using static routing, the most cost-effective method is to use service chaining with UDRs.
This allows:
* Traffic from on-premises to traverse the VPN gateway
* Routing across peered VNets without deploying additional gateways
* No additional ExpressRoute or firewall costs
Microsoft documentation states:
"User-defined routes enable custom routing across peered virtual networks for transit scenarios." Azure Firewall and Application Gateway add unnecessary cost. ExpressRoute is the most expensive option.

NEW QUESTION: 102
Azure □□□□ □□□□ □□□ □□ □ □□ □□ □□ □□□ □□□□ □□□□.
□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.

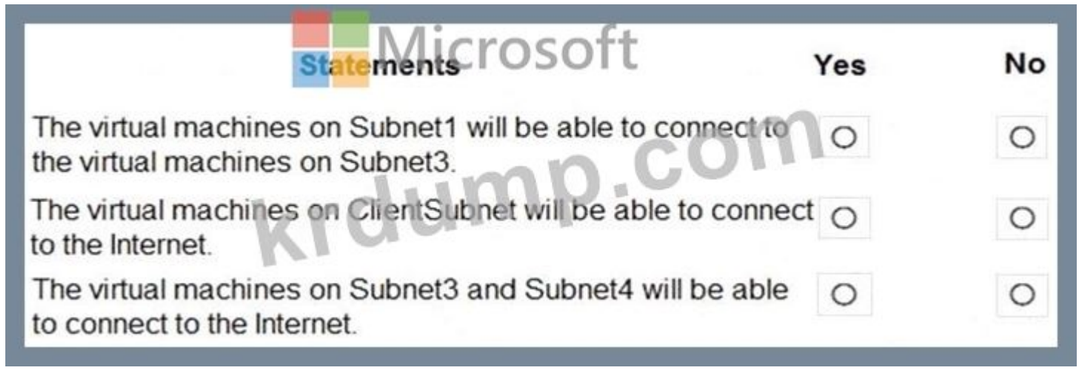

Answer:

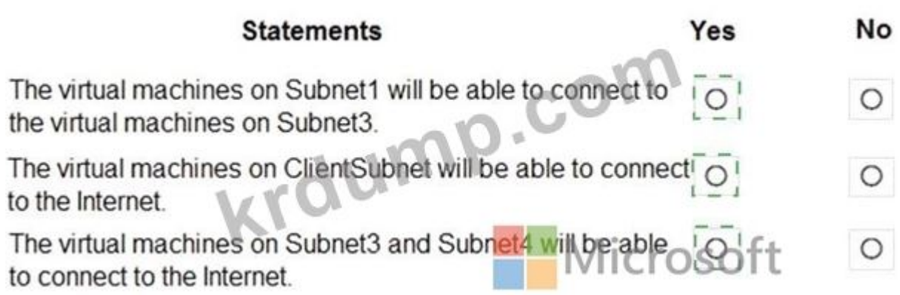

Explanation:

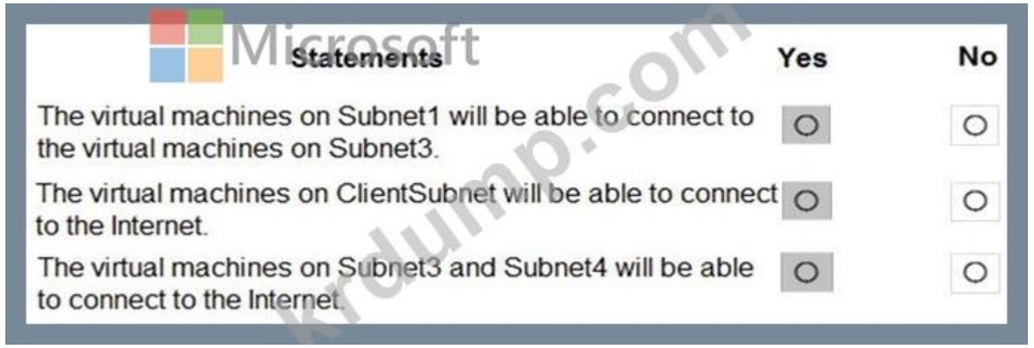

Once the VNets are peered, all resources on one VNet can communicate with resources on the other peered VNets. You plan to enable peering between Paris-VNet and AllOffices-VNet. Therefore VMs on Subnet1, which is on Paris-VNet and VMs on Subnet3, which is on AllOffices-VNet will be able to connect to each other.

All Azure resources connected to a VNet have outbound connectivity to the Internet by default. Therefore VMs on ClientSubnet, which is on ClientResources-VNet will have access to the Internet; and VMs on Subnet3 and Subnet4, which are on AllOffices-VNet will have access to the Internet.

Reference:

https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-peering-overview

https://docs.microsoft.com/en-us/azure/networking/networking-overview#internet-connectivity

Topic 5, Litware, inc.Overview

Litware, Ltd. is a consulting company that has a main office in Montreal and two branch offices in Seattle and New York.

The Montreal office has 2,000 employees. The Seattle office has 1,000 employees. The New York office has

200 employees.

All the resources used by Litware are hosted on-premises.
Litware creates a new Azure subscription. The Azure Active Directory (Azure AD) tenant uses a domain named Litware.onmicrosoft.com. The tenant uses the P1 pricing tier.
Existing Environment
The network contains an Active Directory forest named Litware.com. All domain controllers are configured as DNS servers and host the Litware.com DNS zone.
Litware has finance, human resources, sales, research, and information technology departments. Each department has an organizational unit (OU) that contains all the accounts of that respective department. All the user accounts have the department attribute set to their respective department. New users are added frequently.
Litware.com contains a user named User1.
All the offices connect by using private links.
Litware has data centers in the Montreal and Seattle offices. Each data center has a firewall that can be configured as a VPN device.
All infrastructure servers are virtualized. The virtualization environment contains the servers in the following table.

Name	Role	Contains virtual machine
Server1	VMWare vCenter server	VM1
Server2	Hyper-V-host	VM2

Litware uses two web applications named App1 and App2. Each instance on each web application requires
1GB of memory.
The Azure subscription contains the resources in the following table.

Name	Type
VNet1	Virtual network
VM3	Virtual machine
VM4	Virtual machine

The network security team implements several network security groups (NSGs).
Planned Changes
Litware plans to implement the following changes:
* Deploy Azure ExpressRoute to the Montreal office.
* Migrate the virtual machines hosted on Server1 and Server2 to Azure.
* Synchronize on-premises Active Directory to Azure Active Directory (Azure AD).
* Migrate App1 and App2 to two Azure web apps named webApp1 and WebApp2.
Technical requirements
Litware must meet the following technical requirements:
* Ensure that WebApp1 can adjust the number of instances automatically based on the load and can scale up to five instance*.
* Ensure that VM3 can establish outbound connections over TCP port 8080 to the applications servers in the Montreal office.
* Ensure that routing information is exchanged automatically between Azure and the routers in the Montreal office.
* Enable Azure Multi-Factor Authentication (MFA) for the users in the finance department only.
* Ensure that webapp2.azurewebsites.net can be accessed by using the name app2.Litware.com.
* Connect the New Your office to VNet1 over the Internet by using an encrypted connection.
* Create a workflow to send an email message when the settings of VM4 are modified.
* Create a custom Azure role named Role1 that is based on the Reader role.
* Minimize costs whenever possible.

NEW QUESTION: 103
Azure □□□ □□□□. □□□□ Windows Server□ □□□□ □□ □□□ □□□□ □□□□.
Rule1□□□ □□□ □□□ □□ □□(DCR)□ □□□□.

Azure Monitor Agent□ □□□□ Windows □□□ □□□ □□□□ □□□□ □□□ □□□□□.
ID□ 1001□ □□□ □□□□ □□□□ □□□.
Rule!□ □□□ □□□ □□ □□ □□□ □□□ □□□□ □□□?
A. SQL
B. KQL
C. XPath
Answer: (SHOW ANSWER)
When collecting Windows Event Logs using the Azure Monitor Agent (AMA) with a Data Collection Rule (DCR), filtering is done at the source using XPath queries. XPath is specifically designed for querying XML- based event log entries, including filtering by Event ID, Event Level, and Event Source.
According to Azure Monitor documentation, Windows Event Log data sources require XPath expressions, not SQL or KQL. KQL is used after ingestion for querying data in Log Analytics, while XPath determines which events are collected in the first place.
Since the requirement is to collect only system events with Event ID 1001, the correct query type for the DCR data source is XPath.
Final Answer: C. XPath

NEW QUESTION: 104
Azure □□□ □□□□.
□□□□□ □□□ □□□□□.
□□□□ □□□ □□□□ □□□ □ □□ Azure □□□□ □□□□ □□□.
□□□ □□□□ □□□?
A. Azure □□□□ □ □□
B. Azure Container Instances□
C. Azure Container Apps □□ Azure App Service □□
D. Azure Container Instances □□ Azure App Service □□
E. Azure □□□□ □, Azure □□□□ □□□□ □□ Azure App Service
Answer: C (LEAVE A REPLY)
According to Microsoft Azure Administrator (AZ-104) study guides and Microsoft's official documentation on container hosting options, the ability to automatically scale containerized workloads depends on the underlying platform's orchestration and scaling capabilities.
Azure Container Apps (ACA) is a fully managed service that runs microservices and containers without requiring orchestration infrastructure such as Kubernetes. It natively supports automatic scaling (autoscaling) through KEDA (Kubernetes Event-Driven Autoscaler) integration. Autoscaling in ACA can be triggered by various metrics-such as HTTP request rate, CPU utilization, memory usage, or custom event sources like Azure Service Bus or Event Hub. Administrators can configure minimum and maximum replicas, and Azure dynamically adjusts container instances to match load demand.
Azure App Service can also host containerized web applications (via custom Docker images) and includes built-in autoscaling capabilities. App Service Plans support both manual and automatic scaling based on metrics like CPU percentage, memory percentage, and HTTP queue length. This scaling can occur horizontally (by adding instances) or vertically (by changing the pricing tier).
In contrast, Azure Container Instances (ACI) provides simple, isolated container execution and does not support automatic scaling. You must manually script or orchestrate scaling logic via Azure Automation or external schedulers.
Therefore, per official Microsoft study content and Azure documentation, only Azure Container Apps and Azure App Service offer native, built-in autoscaling for containers.
Final Verified Answer: # C. Azure Container Apps or Azure App Service only

NEW QUESTION: 105
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.

10□□ □□ □□□□□ □□□□ Azure □□□ □□□□. □□ □□□□□ □□□ □□□ □□□□ □□□□□□.
□□ □□□□ □□□ □□ NSG(□□□□ □□ □□)□ □□ □□□□□. NSG□ □□□ □ □□ □□□□ □□ TCP □□ 8080□ □□□□ □□□□□ □□□□ □□□.
□□ □□: □□□ □□ □□ □□ □□□ □□□□□.
□□□ □□□ □□□□□?

A. □

B. □□□

Answer: B (LEAVE A REPLY)

No, this does not meet the goal. Assigning a built-in policy definition to the subscription is not enough to ensure that when an NSG is created, it automatically blocks TCP port 8080 between the virtual networks. This is because there is no built-in policy definition that matches this requirement. The closest built-in policy definition is "Network security groups should not allow unrestricted inbound traffic on well-known ports", but this policy only blocks TCP port 80 and 443, not 80801.

To meet the goal, you need to create a custom policy definition that enforces a default security rule for NSGs. A policy definition is a set of rules and actions that Azure performs when evaluating your resources2.

You can use a policy definition to specify the required properties and values for NSGs, such as the direction, protocol, source, destination, and port of the security rule. You can then assign the policy definition to the subscription scope, so that it applies to all the resource groups and virtual networks in the subscription.

NEW QUESTION: 106

□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□ □□ □□□ □□ □□□ □□□□ □□□□.
VM1□ VM2□□ □ □□ Azure □□ □□□ □□□ App1□□□ □□ □□□□.
Appl□ □□ □□□ Azure Load Balancer□ □□□□ □□□□□.
VM2□ □□ □□□□ □□□□ □□ □□□ □□ □□ □□ □□□□.

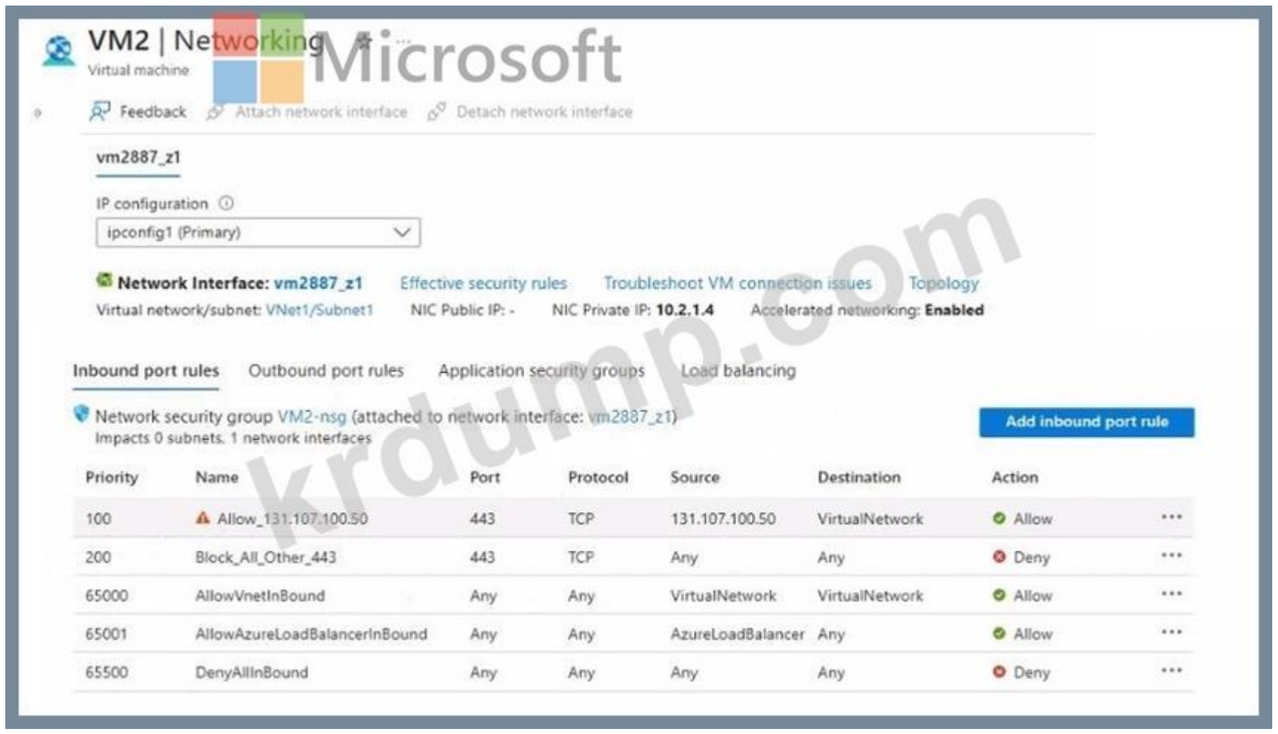

131.107.100.50□□ TCP □□ 443□ □□ 10 Appl □□□ □□□□ □□ □□□□□□.

□□ □□□ □□□ □□□□ □□□□□□ □□□□□.
131.107.100.50□□ TCP □□ 443□ □□ Appl□ □□□□□ □□□ □ □□□ □□□□ □□□.
□□ □□: Azureload Balancer □□□□ □□ □□□□ □□□□ □□□□□ 150□□ □□□□ □□□□ □□ □□□ □□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: (SHOW ANSWER)
In this scenario, the problem involves a Network Security Group (NSG) configuration that determines inbound traffic rules for a virtual machine. The NSG rules for VM2 are evaluated in order of priority, from the lowest number (highest priority) to the highest number (lowest priority).
According to the exhibit, the effective NSG inbound rules for VM2 are:
Priority
Name
Port
Protocol
Source
Destination
Action
100
Allow_131.107.100.50
443
TCP
131.107.100.50
VirtualNetwork
Allow
200
Block_All_Other_443
443
TCP
Any
Any
Deny
65000
AllowVNetInBound
Any
Any
VirtualNetwork
VirtualNetwork
Allow
65001
AllowAzureLoadBalancerInBound
Any

Any
AzureLoadBalancer
Any
Allow
65500
DenyAllInBound
Any
Any
Any
Any
Deny
Root Cause Analysis:
You discovered that connections to App1 from 131.107.100.50 over TCP port 443 fail, even though the Load Balancer and backend pool are configured correctly.
From the rules shown, rule 100 (Allow_131.107.100.50) should theoretically allow that specific IP to connect on port 443. However, note that the Destination for the rule is VirtualNetwork, meaning that the traffic from 131.107.100.50 must be within the same virtual network address space to be permitted.
Since 131.107.100.50 is a public IP address (external source) and not part of the Virtual Network (VNet) address space, the rule does not apply, and the connection fails.
Additionally, the next rule (200 - Block_All_Other_443) explicitly denies any other inbound traffic on TCP
443 from any source. Therefore, the inbound traffic from 131.107.100.50 is blocked.
Evaluation of the Proposed Solution:
The proposed solution suggests creating an inbound security rule:
"Allow any traffic from the AzureLoadBalancer source and set priority to 150." However, the existing NSG already includes a built-in rule (65001 - AllowAzureLoadBalancerInBound) that allows inbound traffic from the Azure Load Balancer. That means traffic originating from the Azure Load Balancer front-end is already allowed by default.
Since the failed connection is coming from an external client (131.107.100.50) - not from the Load Balancer source itself - adding another rule allowing AzureLoadBalancer traffic does not resolve the issue. The correct action would be to modify or create a new rule that:
Allows inbound traffic on port 443
From source 131.107.100.50
With destination = Any
Priority lower than 200 (i.e., higher precedence than the deny rule)
Verified Microsoft Azure Administrator Reference:
From Microsoft Docs - Network security group overview and priority order:
"Azure processes the security rules in priority order, starting from the lowest number. Once a rule matches the traffic, processing stops. Lower-numbered priority rules override higher-numbered ones."
"The built-in rule AllowAzureLoadBalancerInBound allows traffic from Azure Load Balancer but not from external public IPs directly accessing the virtual machine." Correct Resolution:
You need to create a new inbound rule allowing TCP 443 from 131.107.100.50 with Destination: Any, Priority: <200, for example:
Priority: 150
Source: IP address (131.107.100.50)
Destination: Any
Protocol: TCP
Port: 443
Action: Allow
This will ensure the traffic is allowed before the deny rule at 200 takes effect.
Final Verified Answer:

B). No
The proposed rule allowing traffic from AzureLoadBalancer does not help, because traffic from
131.107.100.50 originates externally, not from the Load Balancer. You must instead create a rule that explicitly allows that IP on port 443 with a higher priority than the existing deny rule.

NEW QUESTION: 107
VNet1□□□ □□ □□□□□ □□ □□ Azure □□ □□□ □□□□.
□□ □□□ □□ Azure Storage □□□ □□□□□.

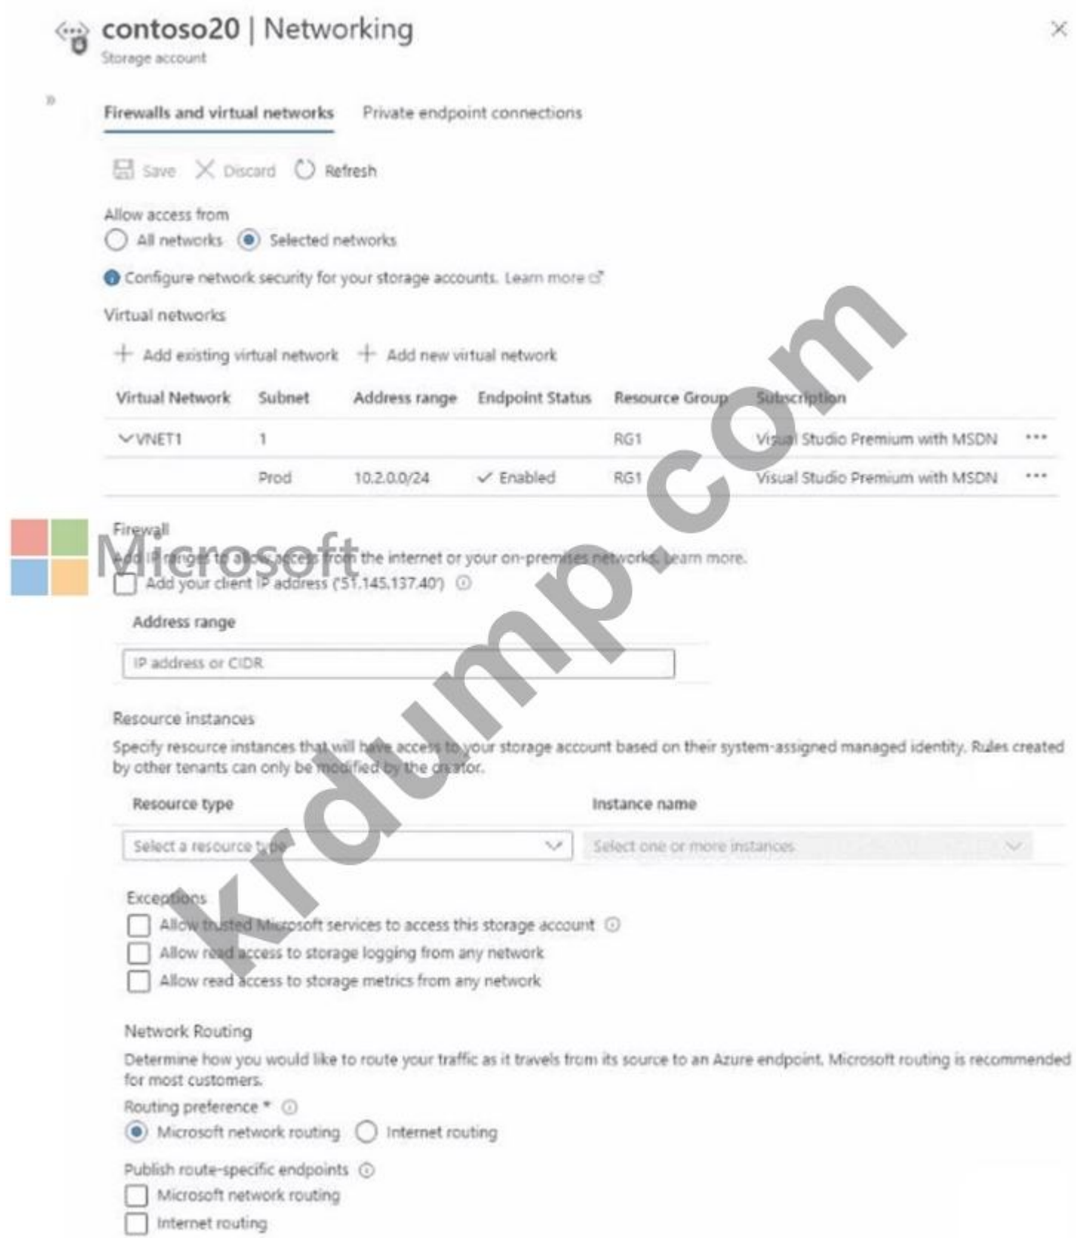
contoso20 | Networking
Storage account
Firewalls and virtual networks
Private endpoint connections
Save
Discard
Refresh
Allow access from
All networks
Selected networks
Configure network security for your storage accounts. Learn more
Virtual networks
Add existing virtual network
Add new virtual network
Virtual Network
Subnet
Address range
Endpoint Status
Resource Group
Subscription
VNET1
1
RG1
Visual Studio Premium with MSDN
Prod
10.2.0.0/24
Enabled
RG1
Visual Studio Premium with MSDN
Firewall
Add your client IP address ('51.145.137.40')
Address range
IP address or CIDR
Resource instances
Specify resource instances that will have access to your storage account based on their system-assigned managed identity. Rules created by other tenants can only be modified by the creator.
Resource type
Instance name
Select a resource type
Select one or more instances
Exceptions
Allow trusted Microsoft services to access this storage account
Allow read access to storage logging from any network
Allow read access to storage metrics from any network
Network Routing
Determine how you would like to route your traffic as it travels from its source to an Azure endpoint. Microsoft routing is recommended for most customers.
Routing preference *
Microsoft network routing
Internet routing
Publish route-specific endpoints
Microsoft network routing
Internet routing

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

Never

Never

NEW QUESTION: 108

Microsoft Entra □□□□ □□□□.

□ Microsoft 365 □□□ □□□□ □□ □□□ □□□ □□ □□□□ □□□ □□□□□ □□□□ □□□.

<□□x□□ □□>.

Microsoft Entra □□ □□□□ □□ □ □□ □□□ □□□□ □□□□ □□□? □□□□□ □□ □□□□ □□ □□□ □□ □□□□ □□□ □□□ □□□□ □□□□□.

Actions

- Customize the company branding.
- Set Add suffix to **String**.
- Set Add suffix to **Attribute**.
- Set Add prefix to **String**.
- Create a group naming policy.
- Set Add prefix to **Attribute**.
- Set Select type to **Department**.

Answer Area

Answer:

Actions

- Customize the company branding.
- Set Add suffix to **String**.
- Set Add suffix to **Attribute**.
- Set Add prefix to **String**.
- Create a group naming policy.
- Set Add prefix to **Attribute**.
- Set Select type to **Department**.

Answer Area

- Create a group naming policy.
- Set Add prefix to **Attribute**.
- Set Select type to **Department**.

Explanation:

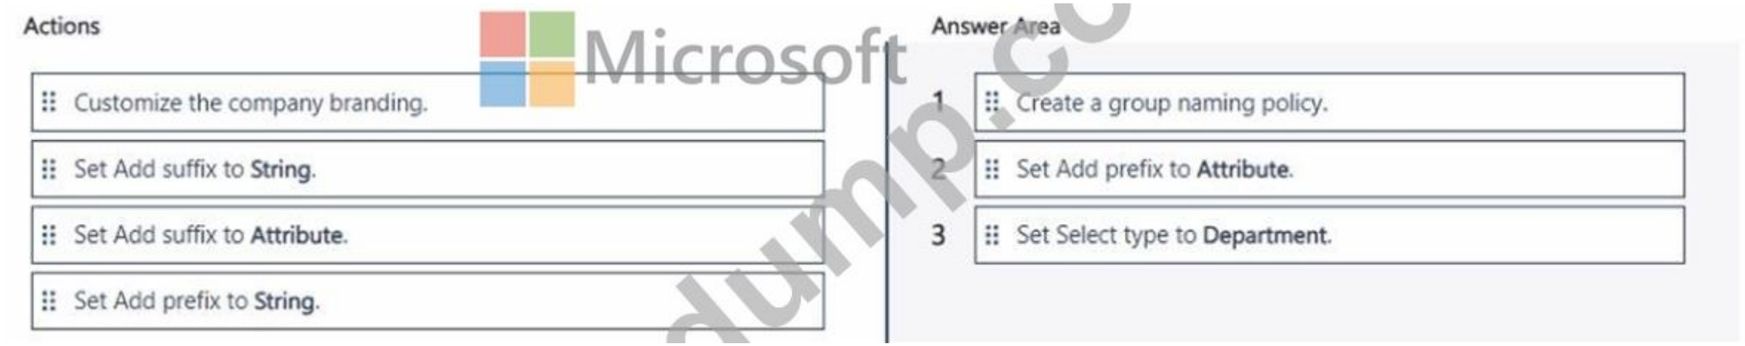

Microsoft Entra ID (formerly Azure Active Directory) supports group naming policies to automatically enforce consistent and compliant naming conventions for Microsoft 365 groups. The purpose of this policy is to standardize names, prevent conflicts, and include metadata such as department, location, or purpose.

According to the Microsoft Entra Administrator documentation and the AZ-104 official study guide ("Manage Azure Identities and Governance"), the steps to implement such a policy are:

Step 1: Create a Group Naming Policy

Navigate to:

Microsoft Entra admin center # Groups # Naming policy.

Here, you can define global naming conventions that apply whenever new Microsoft 365 groups or security groups are created.

Creating the policy allows you to specify prefixes and suffixes that automatically appear in group names based on attributes or fixed strings.

Step 2: Set Add Prefix to Attribute

Once the naming policy is created, select the Add prefix option.

You can choose between two prefix types:

String: A fixed word or phrase (for example, "Corp_")

Attribute: A user or group attribute dynamically pulled from Microsoft Entra ID (for example, Department or CompanyName).

To achieve the naming format <Department><Group name>, you must choose Attribute as the prefix type, since it will automatically insert the department value of the creator or group owner.

Step 3: Set Select Type to Department

After choosing Attribute as the prefix, you select which attribute will be used.

Choose Department to ensure the prefix dynamically reflects the department name of the group creator.

This results in an automatically generated group name like:

FinanceSales, HRRecruiting, etc.

If desired, you could later add a suffix (like "_Group"), but for this specific scenario, only the prefix (Department) is required.

Official Microsoft Documentation Extract (Summarized):

"A naming policy can consist of prefixes or suffixes that include fixed strings or user attributes, such as

[Department], [Company], or [Office].

To configure, create a group naming policy and add prefix or suffix elements based on attributes." (Reference: Microsoft Learn - Configure naming policy for Microsoft 365 groups in Azure Active Directory)

NEW QUESTION: 109

Subscription1□□□ Azure □□□ □□□□.

Subscription1□□ share1□□□ Azure □□ □□□ □□□□.

□□ □□□ □□ SAS1□□□ □□□ □□ □□□ □□(SAS)□ □□□□.

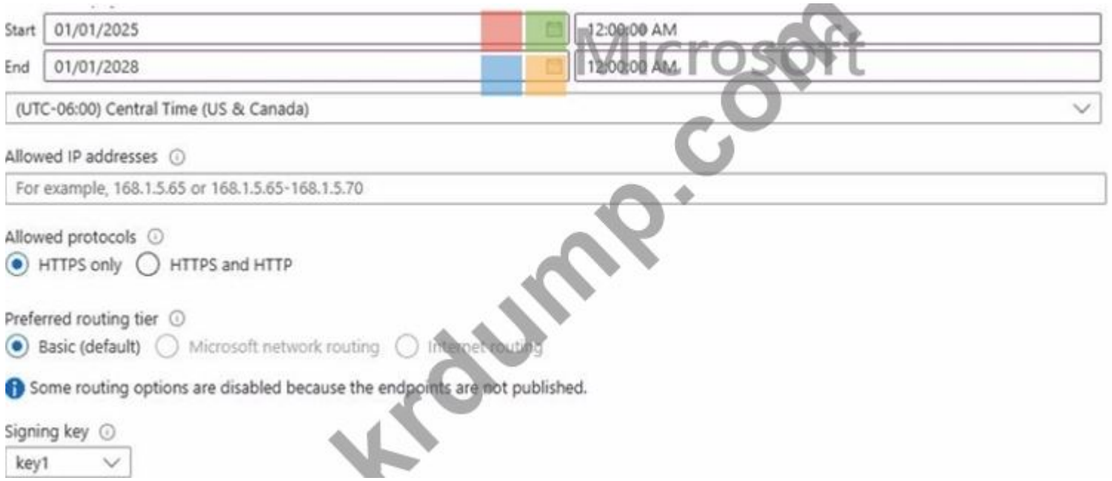
Start 01/01/2025 12:00:00 AM
End 01/01/2028 12:00:00 AM
(UTC-06:00) Central Time (US & Canada)
Allowed IP addresses
For example, 168.1.5.65 or 168.1.5.65-168.1.5.70
Allowed protocols
HTTPS only
HTTPS and HTTP
Preferred routing tier
Basic (default)
Microsoft network routing
Internet routing
Some routing options are disabled because the endpoints are not published.
Signing key
key1

□□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.


Answer Area
If on January 2, 2025, you run Microsoft Azure Storage Explorer on a computer that has an IP address of 193.77.134.1, and you use SAS1 to connect to the storage account, you [answer choice].
will be prompted for credentials
Will have no access
will have read, write, and list access
will have read-only access
If on January 10, 2025, you run the net use command on a computer that has an IP address of 193.77.134.50, and you use SAS1 as the password to connect to share1, you [answer choice].
will be prompted for credentials
will have no access
Will have read, write, and list access
will have read-only access

Answer:

Explanation:

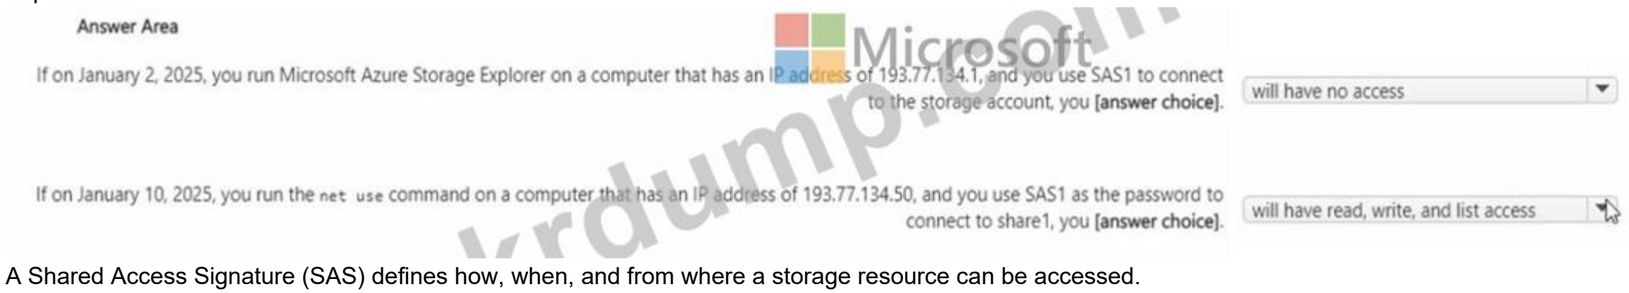

A Shared Access Signature (SAS) defines how, when, and from where a storage resource can be accessed.
The exhibit shows the following important configuration details for SAS1:
Validity period: January 1, 2025 through January 1, 2028 # both access attempts occur within the valid timeframe.
Allowed IP addresses: Not specified # access is allowed from any IP address.
Allowed protocols: HTTPS only
Signing key: key1
Scenario 1 - Azure Storage Explorer
Azure Storage Explorer accesses Azure Storage using the HTTPS REST API. Because:
The access is within the SAS validity period,
No IP restriction is configured,
HTTPS is allowed,
the connection succeeds. However, the SAS configuration shown does not include write or list permissions, meaning access is limited to read-only operations.
Result: Will have read-only access
Scenario 2 - net use to connect to Azure file share
The net use command connects to Azure Files using SMB (port 445). SMB traffic does not use HTTPS.
Because the SAS explicitly restricts access to HTTPS only, SMB-based access is blocked, regardless of IP address or validity period.
Microsoft Azure documentation clearly states:
"When HTTPS only is selected, requests that use HTTP or SMB are denied."
Result: Will have no access

Final Answer Summary
Scenario
Result
Storage Explorer over HTTPS
Will have read-only access
SMB (net use) connection
Will have no access
Final Verified Answer:
First scenario: Will have read-only access
Second scenario: Will have no access

NEW QUESTION: 110
□□ □□ □□ □□ □□□□ □□□ contoso2024□□ Azure Storage □□□ □□□□.

Name	Type	Contents
container1	Blob container	File1
share1	Azure Files share	File2

□□ □□ □□□ □□□ contoso2024□ □□ □□□ □□ □□□□ □□□□.

Name	Permission
User1	Reader role
User2	Storage Account Contributor role
User3	Has an access key for contoso2024

contoso2024 □□□ □□ □□□ □□ □□□□□.

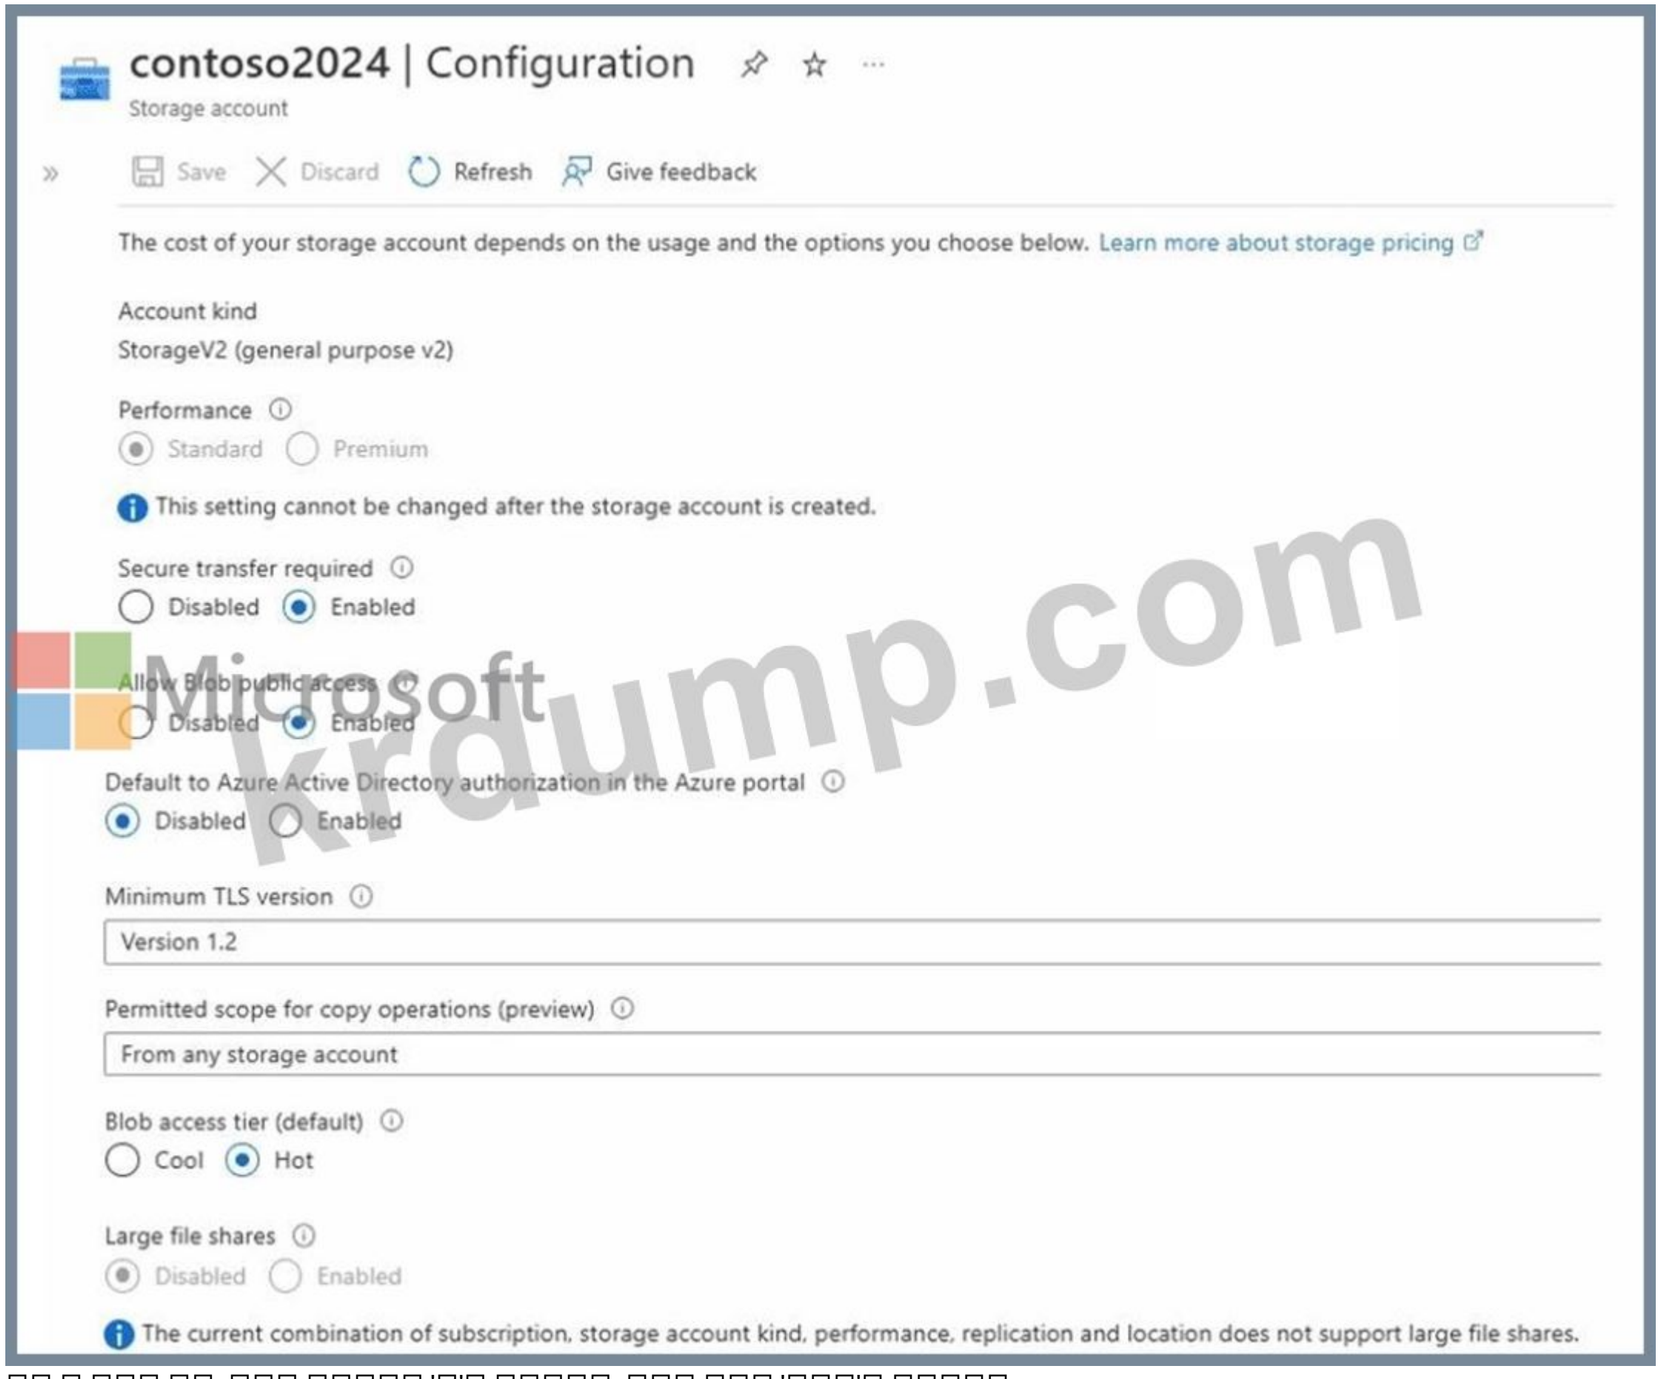
contoso2024 | Configuration
Storage account
Save
Discard
Refresh
Give feedback
The cost of your storage account depends on the usage and the options you choose below. Learn more about storage pricing
Account kind
StorageV2 (general purpose v2)
Performance
Standard
Premium
This setting cannot be changed after the storage account is created.
Secure transfer required
Disabled
Enabled
Allow Blob public access
Disabled
Enabled
Default to Azure Active Directory authorization in the Azure portal
Disabled
Enabled
Minimum TLS version
Version 1.2
Permitted scope for copy operations (preview)
From any storage account
Blob access tier (default)
Cool
Hot
Large file shares
Disabled
Enabled
The current combination of subscription, storage account kind, performance, replication and location does not support large file shares.

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
User1 can read File1.	○	○
User2 can read File2.	○	○
User3 can read File1 and File2.	○	○

Answer:

Answer Area

Statements	Yes	No
User1 can read File1.	◉	○
User2 can read File2.	○	◉
User3 can read File1 and File2.	○	◉

NEW QUESTION: 111

□□ 1□□□ Azure □□□ □□ Microsoft System Center Service Manager □□□ □-□□□□ □□□□ VM1□□□ □□ □□□ □□□□ □□□□.

VM1□ □□ □□□ □□□ □□ 10% □□□ □ □□□ □□□□ □□□ □□□□□ □□□□ □□□. □□ □□ □□□ □□ □□□?

A. □□□ □□□□.

B. □□□ □□□ □□□□.

C. IT □□□ □□ □□□(ITSM)□ □□□□□.

D. □□ □ □□

Answer: C (LEAVE A REPLY)

IT Service Management Connector (ITSMC) allows you to connect Azure to a supported IT Service Management (ITSM) product or service. Azure services like Azure Log Analytics and Azure Monitor provide tools to detect, analyze, and troubleshoot problems with your Azure and non-Azure resources. But the work items related to an issue typically reside in an ITSM product or service. ITSMC provides a bi-directional connection between Azure and ITSM tools to help you resolve issues faster. ITSMC supports connections with the following ITSM tools: ServiceNow, System Center Service Manager, Provance, Cherwell.

Reference:

https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/itsmc-overview

NEW QUESTION: 112

□□ □□ □□□ □□□□ □□□□ Subscription1□□□ Azure □□□ □□□□.

Quota name	Region	Current Usage
Standard BS Family vCPUs	West US	0 of 20
Standard D Family vCPUs	West US	0 of 20
Total Regional vCPUs	West US	0 of 20

□□ □□ □□□ □□ Subscription1□ □□ □□□ □□□□□.

Name	Size	vCPUs	Region	Status
VM1	Standard_B2ms	2	West US	Running
VM2	Standard_B16ms	16	West US	Stopped (Deallocated)

□□ □□ □□□ □□ □□□ □□□ □□□□□.

Name	Size	vCPUs
VM3	Standard_B2ms	2
VM4	Standard_D4s_v3	4

□□ □ □□□ □□, □□□ □□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□. □□: □□□ 1□□□□.

Answer:

Answer Area

Explanation:
VM3 - Yes
VM4 - No
VM5 - No

NEW QUESTION: 113
□ □□ □□ □□ □□□□□ □□□□ App1□□□ Azure App Service □□ □□□□.
□□ □□□ □□ □□ □□ □□ □□□ □□□□ □□□□.

Criteria

Metric namespace *

Standard metrics

Metric name

Memory Percentage

1 minute time grain

Dimension Name	Operator	Dimension Values	Add
Instance	=	All values	+

If you select multiple values for a dimension, autoscale will aggregate the metric across the selected values, not evaluate the metric for each values individually.

MemoryPercentage (Average)

39.28 %

Enable metric divide by instance count

Operator *

Greater than

Metric threshold to trigger scale action *

70

%

Duration (minutes) *

15

Time grain (minutes)

1

Time grain statistic *

Average

Time aggregation *

Average

Action

Operation *

Increase count by

Cool down (minutes) *

5

instance count *

1

□□□□ □□ □□ □□ □□□ □□ □□□□ 5□ □□□□□.
30□ □□ App1□ □□ □□□ □□□□ 80%□ □□□□□.
30□ □□ App1□ □□ □□□□ □□ □ □□□□?
A. 2
B. 3
C. 4
D. 5
Answer: (SHOW ANSWER)
In Azure App Service, autoscale rules automatically adjust the number of running instances based on performance metrics such as CPU, memory, or custom metrics. The configuration shown uses the Memory Percentage metric with a threshold of 70%, meaning that if average memory utilization exceeds 70% for 15 minutes, Azure will trigger a scale-out action.
From the exhibit, the rule specifies:
* Metric threshold: Greater than 70%
* Duration: 15 minutes
* Action: Increase count by 1 instance
* Cool down: 5 minutes
* Maximum instance limit: 5
App1 currently has 2 running instances.
If App1 maintains 80% memory utilization for 30 minutes, the autoscale mechanism will trigger the scale-out action after each qualifying 15-minute window, adding 1 instance every 5-minute cooldown period until the maximum instance limit (5) is reached.
Therefore, the progression would be:
* Start: 2 instances
* After first 15 minutes: +1 instance # 3 total
* After next 5-minute cooldown and continuing high memory # +1 instance # 4 total
* After another cycle # +1 instance # 5 total
Since the rule continues to trigger until the maximum (5) is reached, and the memory usage remains above the threshold, the maximum number of instances that App1 will scale to during the 30-minute period is 5.

NEW QUESTION: 114
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.
VM1□□□ Azure □□ □□□ □□□□. VM1□ ARM1.json□□□ □□□ □□ Azure Resource Manager □□□□ □□□□ □□□□□□□.
VM1□ □□ □□□ □□□ □□ □□□□ □□□ □□□□.
VM1□ □□ □□ □□□□ □□□□ □□□.
□□ □□: □□ □□□□□□ □□ □□□ □□ □□□ □□□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
Moving the virtual machine to a different resource group does not change the host that the virtual machine runs on. It only changes the logical grouping of the resources. To move the virtual machine to a different host, you need to redeploy it or use Azure Site Recovery. Then, References: [Move resources to new resource group or subscription] [Redeploy Windows VM to new Azure node] [Use Azure Site Recovery to migrate Azure VMs between Azure regions]

NEW QUESTION: 115
Azure □□□ □□□□.
File1.bicep□□□ □□□ □□ □□□ □□□ □□□□□.

```
param location string = resourceGroup().location

resource virtualNetwork 'Microsoft.Network/virtualNetworks@2024-01-01' = {
  name: 'VNET1'
  location: location
  properties: {
    addressSpace: {
```

Answer:

Explanation:

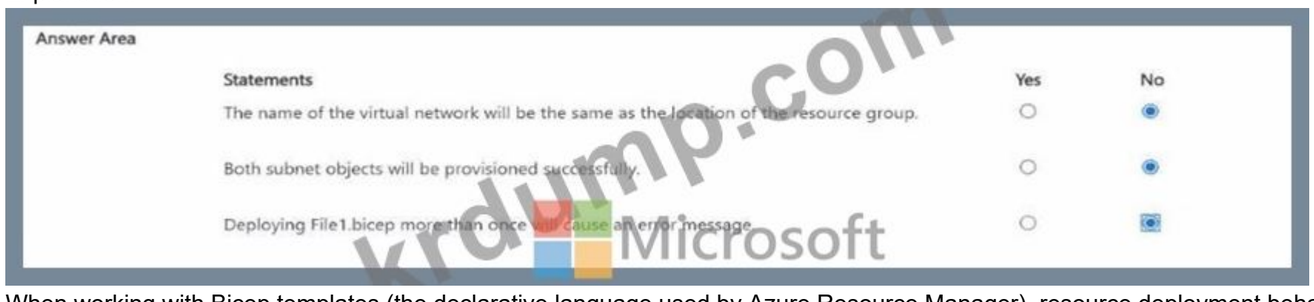

When working with Bicep templates (the declarative language used by Azure Resource Manager), resource deployment behaviors are governed by parameters, variable assignments, and idempotent deployment principles.
Name of the Virtual Network:
In a Bicep file, the name of a virtual network (VNet) is explicitly defined in the template or passed as a parameter. It does not automatically inherit the name or location of the resource group. Therefore, unless explicitly coded that way, the VNet name is not derived from the resource group's location or name.
Subnet provisioning:
A virtual network must have unique subnet names and non-overlapping address spaces. If a Bicep template contains subnet definitions with conflicting address prefixes or duplicated names, Azure Resource Manager will fail to provision them. Thus, if both subnets have overlapping ranges or identical names, both will not be provisioned successfully.

Re-deployment of File1.bicep:
Azure Bicep and ARM templates are idempotent-meaning repeated deployments of the same template with the same parameters do not cause errors. Instead, existing resources are verified and updated if necessary.
Hence, redeploying File1.bicep will not cause an error message unless a conflict or manual deletion occurs.

NEW QUESTION: 116
East US 2 □□□ VNET1□□□ □□ □□□□□ □□□ Azure □□□ □□□□.
VM1-NI □ VM2-NI□□ □□□□ □□□□□□ VNET1□ □□□□□.

```
{
    "apiVersion": "2024-07-01",
    "type": "Microsoft.Compute/virtualMachines",
    "name": "VM1",
    "zones": "1",
    "location": "EastUS2",
    "dependsOn": [
      "[resourceId('Microsoft.Network/networkInterfaces', 'VM1-NI')]"
    ],
    "properties": {
      "hardwareProfile": {
        "vmSize": "Standard_A2_v2"
      },
      "osProfile": {
        "computerName": "VM1",
        "adminUsername": "AzureAdmin",
        "adminPassword": "[parameters('adminPassword')]"
      },
      "storageProfile": {
        "osDisk": {
          "createOption": "FromImage"
        }
      },
      "networkProfile": {
        "networkInterfaces": [
          {
            "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM1-NI')]"
          }
        ]
      }
    }
},
{
    "apiVersion": "2024-07-01",
    "type": "Microsoft.Compute/virtualMachines",
    "name": "VM2",
    "zones": "2",
    "location": "EastUS2",
    "dependsOn": [
      "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
    ],
    "properties": {
      "computerName": "VM2",
      "adminUsername": "AzureAdmin",
      "adminPassword": "[parameters('adminPassword')]"
    },
    "storageProfile": {
      "imageReference": "[variables('image')]",
      "osDisk": {
        "createOption": "FromImage"
      }
```

```
    },
    "networkProfile": {
      "networkInterfaces": [
        {
          "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
        }
      ]
    }
  }
}
```

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

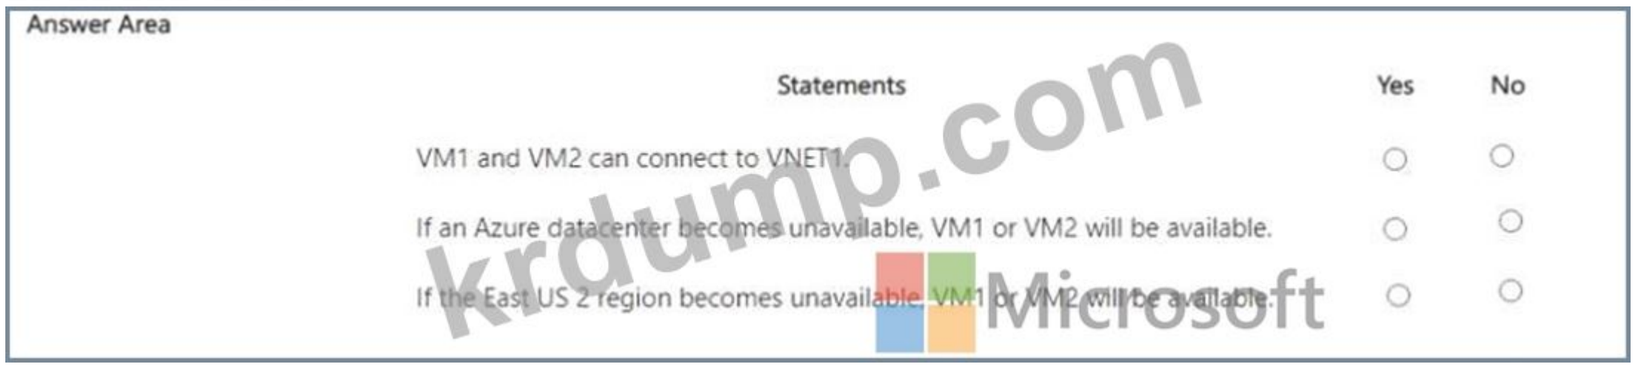

Answer:

Explanation:

The ARM template configuration shows that both VM1 and VM2 are deployed to East US 2 and each VM references its respective network interface in the networkProfile by resource ID. Because the prompt states that VM1-NI and VM2-NI are connected to VNET1, both VMs are therefore attached to VNET1 through their NICs. In Azure, a VM connects to a virtual network via the NIC resource; when the NIC is attached to a subnet in the VNet, the VM has network connectivity to that VNet.

Answer Area

The template also sets availability zones using the zones property: VM1 is deployed to zone 1 and VM2 to zone 2 within the same region. Microsoft Azure documentation for Availability Zones explains that zones are physically separate locations (datacenters) within an Azure region, each with independent power, cooling, and networking. Therefore, if a single datacenter/zone becomes unavailable, workloads placed in a different zone can remain available, meaning at least one of the VMs (VM1 or VM2) should still be available.
However, Availability Zones do not protect against a regional outage. Since both VMs are deployed to East US 2, if the entire region becomes unavailable, neither VM would be available without a multi-region architecture (for example, paired-region deployment, cross-region replication, or failover).

NEW QUESTION: 117
storageaccount1□□□ Azure Storage □□□ □□□ Azure □□□ □□□□.
storageaccount1□ Azure Resource Manager □□□□□ □□□□□. □□□□□ □□ □□□ □□□□ □□□□.

```
{
        "type": "Microsoft.Storage/storageAccounts",
        "apiVersion": "2019-06-01",
        "name": "storageaccount1",
        "location": "eastus",
        "sku": {
            "name": "Standard_LRS",
            "tier": "Standard"
        },
        "kind": "StorageV2",
        "properties": {
            "networkAcls": {
                "bypass": "AzureServices",
                "virtualNetworkRules": [],
                "ipRules": [],
                "defaultAction": "Allow"
            },
            "supportsHttpsTrafficOnly": true,
            "encryption": {
                "services": {
                    "file": {
                        "keyType": "Account",
                        "enabled": true
                    },
                    "blob": {
                        "keyType": "Account",
                        "enabled": true
                    }
                },
                "keySource": "Microsoft.Storage"
            },
            "accessTier": "Hot"
        }
    },
```

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer:

Explanation:

The provided ARM template defines a storage account named storageaccount1 of kind StorageV2 and SKU Standard_LRS.

1## Public IP Access (Yes):

In the template, the networkAcls property sets "defaultAction": "Allow", meaning that access to the storage account is open to all public networks by default. Since no ipRules or virtualNetworkRules are defined, there are no restrictions based on IP addresses. Therefore, any server - including one with public IP

131.107.103.10 - can connect successfully.

(Reference: Azure Storage firewall and virtual network documentation - when defaultAction is Allow, storage accepts all incoming requests over the public endpoint.)
2## Archive Tier Availability (Yes):
Because the storage account is of kind StorageV2, it supports blob access tiers - Hot, Cool, and Archive - at both the account and individual blob levels. The ARM template specifies "accessTier": "Hot", which is the default access tier for the account. However, this does not restrict individual blobs from being set to a different tier (e.g., Archive) later. Azure Blob Storage in a StorageV2 account allows per-blob tiering, enabling you to optimize costs based on access patterns.
(Reference: Azure Storage documentation - "Blob storage lifecycle management" and "Blob access tiers overview.")
3## Azure AD Authentication for File Shares (No):
The ARM template enables encryption and HTTPS but does not configure Azure Active Directory Domain Services (Azure AD DS) authentication for Azure Files. Global administrators in Azure AD cannot access SMB file shares using their Azure AD credentials unless Azure Files AD DS integration is explicitly configured. Since no such property (e.g., azureFilesIdentityBasedAuthentication) exists in this template, access would still require storage account keys or shared access signatures (SAS), not Azure AD credentials.
(Reference: Azure Files authentication overview - Azure AD DS authentication requires explicit configuration.)
Final Verified Answers:
* A server that has public IP 131.107.103.10 can access storageaccount1 # Yes
* Individual blobs in storageaccount1 can be set to use the archive tier # Yes
* Global administrators in Azure AD can access a file share hosted in storageaccount1 by using their Azure AD credentials # No

NEW QUESTION: 118
contoso.com□□□ □□□ DNS □□□□ □□□□.
contoso.com□□□ □□□ □□ Azure DNS □□□ □□□□.
contoso.com □□□□ □□□ □□□□ □□□□□ □□□ □ □□□ □□□□ □□□.
□□□ □□ □□□?
A. DNS □□□ □□ □□□ NS □□□□ □□□□□.
B. contoso.com□ NS □□□□ □□□□.
C. contoso.com□ SOA □□□□ □□□□.
D. DNS □□□ □□ □□□ SOA □□□□ □□□□□.
Answer: A (LEAVE A REPLY)
When a public Azure DNS zone is created, Azure automatically assigns a set of authoritative name servers (NS records) to the zone. However, simply creating the DNS zone in Azure does not make it resolvable from the internet. For external resolution to work, the domain registrar must delegate authority to Azure DNS.
Microsoft Azure documentation clearly states that to make DNS records resolvable publicly, you must update the domain registrar's NS records to point to the Azure-assigned name servers. This action establishes Azure DNS as the authoritative DNS provider for the domain.
Creating NS or SOA records within the Azure DNS zone itself does not affect external resolution unless the registrar delegation is completed. The SOA record is automatically created and managed by Azure DNS and must not be modified at the registrar.

NEW QUESTION: 119
Admin1□ □□ □□□ □□ □□□ □□□□ □□□.
□□□ □□ □□□?
A. □□ □□□□□□ □□□ □□□ □□ □□□ □□□□□.
B. □□ □□□□□□ □□□ □□□ □□ □□□ □□(IAM) □□□ □□□□□.
C. Azure Active Directory □□□□□□ □□□ □□□□□.
D. Azure Active Directory □□□□□□ □□□ □□□□□.
Answer: A (LEAVE A REPLY)

In this scenario, Contoso Ltd. must designate a new user named Admin1 as the service administrator of the Azure subscription and ensure that Admin1 receives email alerts about service outages.
In Azure, there are three classic administrative roles associated with subscriptions:
* Account Administrator - The individual who created the subscription and manages billing.
* Service Administrator - The primary contact responsible for managing services and resources in the subscription.
* Co-Administrator - Has the same management privileges as the Service Administrator but cannot change subscription associations.
According to Microsoft Azure Administrator documentation, to change the Service Administrator, you must perform the following steps:
* In the Azure portal, navigate to Subscriptions.
* Select the specific subscription.
* Under Settings, choose Properties.
* In the Service Administrator field, update the name and email address to that of the new administrator (in this case, Admin1).
This modification ensures that Admin1 becomes the Service Administrator, and Azure will automatically send service-related notifications and outage alerts to that user's registered email address, as defined by Microsoft' s subscription notification process.
Option B (IAM settings) is used for Role-Based Access Control (RBAC) and assigning Azure Resource Manager roles such as Owner, Contributor, or Reader. However, RBAC roles do not change the Service Administrator at the subscription level - that's only done through the Properties blade.
Options C (AAD Properties) and D (AAD Groups) are unrelated to subscription-level administrative settings.
Therefore, the verified and Microsoft-documented answer is:
A. From the Subscriptions blade, select the subscription, and then modify the Properties.
Final Verified Answer: # A. From the Subscriptions blade, select the subscription, and then modify the Properties.

NEW QUESTION: 120
Microsoft Entra □□□□ □□□□.
Azure Portal□ □□□□ □□□□ □□□□ □□□□□ □□□.
□□□ □□□ □□□ □ □□□ □□□ □□□□ □□ □□□ □□□□□ □□ □□□□□ □□ □□□. □□□□ □□ □□□ □□□□□ □□□.
□□ □ □□ □□□ □□□ □□□? □ □□□ □□□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.
A. Azure Resource Manager(ARM) □□□□ □□□□.
B. □□ □□□ □□□ □□□ □□□□ □□□ □□□□.
C. □□□□ □□□ □□ □□□□ PowerShell □□□□□ □□□□□.
D. □□□ □□□ □□□ □□□□ □□□ □□□□.
E. □□□ □□□ □□□ □□□ □□□ CSV □□□ □□□□.
F. □□□ □□□ □□□ □□□ □□□□ XML □□□ □□□□.
Answer: B,E (LEAVE A REPLY)
Microsoft Entra ID supports bulk user creation through the Azure portal by uploading a CSV file that contains user attributes such as UserPrincipalName, DisplayName, Department, and JobTitle. This method is explicitly documented as the supported approach for bulk importing users and minimizes administrative effort compared to scripting or manual creation.
To ensure that users are automatically added to a group based on their department, Microsoft Entra ID provides Dynamic User groups. Dynamic User membership allows administrators to define membership rules that evaluate user attributes (for example, user.department -eq "Finance"). Once users are imported with the correct department attribute, Entra ID automatically evaluates the rules and adds users to the appropriate group without manual intervention.
Assigned membership groups require administrators to manually add users, and PowerShell scripting increases complexity and maintenance overhead. XML files are not supported for bulk user import in the Azure portal.
By combining Dynamic User membership groups with a CSV import containing department attributes, the solution ensures automation, accuracy, and minimal administrative effort, exactly as required.
Final Answers:
* B. Create groups that use the Dynamic User membership type

* E. Create a CSV file that contains user information and the appropriate attributes

NEW QUESTION: 121

□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Location	IP address space	Subnet
VNet1	East US	10.1.128.0/23	Subnet1
VNet2	East US	192.168.0.0/16	Subnet21, Subnet22
VNet3	East US	172.16.0.0/16	Subnet3

□□ □□□ □□□□ IP □□ □□□ □□□□ □□□□.

Name	IP address space
Subnet1	10.1.128.0/24
Subnet21	192.168.0.0/17
Subnet22	192.168.128.0/17
Subnet3	172.16.1.0/24

□□ □□ Azure □□□ contapp1□□□ □□□□ □□ □□ □□□□□.

□□ □□ □□□ □□□□ con-env1□□□ □□□□ □ □□□ □□□□ □□□.

* □□ □□ □□□□□ □□□□□.
* □□ □□□□ □□□□□.
* □□□ □□ □□ □□□□ □□□□□.

con-env1□ □□ □□ □□□□□ □□□ □ □□□, □□ □□□ □□□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer Area

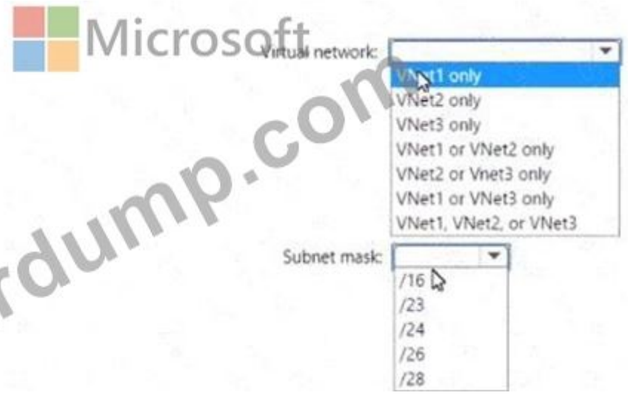

Answer:

Answer Area

Explanation:

Virtual Network: You can connect con-env1 to VNet2 and VNet3 only. This is because VNet1 is in a different region than the container app, which is East US. According to the web search results, you can only connect a container app environment to a virtual network that is in the same region as the container app1.

Therefore, VNet1 is not a valid option. VNet2 and VNet3 are both in the same region as the container app, and they have enough available IP addresses to support a container app environment.

Subnet mask: You should use /28 as the subnet mask for con-env1. This is because /28 is the smallest possible subnet mask that can accommodate a container app environment. According to the web search results, a container app environment requires a minimum of 16 IP addresses in a subnet2. A /28 subnet mask provides

16 IP addresses, while a /26 subnet mask provides 64 IP addresses, a /24 subnet mask provides 256 IP addresses, a /23 subnet mask provides 512 IP addresses, and a /16 subnet mask provides 65,536 IP addresses.

Therefore, /28 is the most efficient choice for minimizing the subnet size.

NEW QUESTION: 122

webapp1□□□ Azure □□□ □□□□.

VNET1□□□ □□ □□□□□ MySQL □□□□□□□ □□□□□ VM1□□□ Azure □□ □□□ □□□□. VM1□ VNET1□ □□□□□.

webapp1□ VM1□ □□□□ □□□□ □□□□ □ □□□ □□□□ □□□.

□□□ □□ □□□?

A. webapp1□ VNET1□ □□□□□.

B. □□ □□ □□□□ □□□□□.

C. Azure Application Gateway □□,

D. VNET1□ □□ □□ □□□□□ □□□□□□.

Answer: (SHOW ANSWER)

Azure App Service (which hosts web apps) can securely connect to resources inside an Azure Virtual Network (VNet) by enabling VNet Integration.

In this scenario:

* VM1 hosts a MySQL database inside VNET1.

* webapp1 (the Azure Web App) needs to access that data privately.

To allow webapp1 to communicate with VM1, you should enable VNet Integration on webapp1, connecting it to VNET1.

This ensures:
* Private, secure communication between the App Service and the virtual machine.
* No need to expose public IPs or deploy extra load balancers.
Incorrect options:
* B. Internal Load Balancer # Used to balance internal traffic among VMs, not for connecting web apps.
* C. Application Gateway # Provides HTTP load balancing, not private connectivity.
* D. VNet Peering # Used to connect different VNets, not to connect an App Service to a VNet.
# Final Verified Answer: A. Connect webapp1 to VNET1

NEW QUESTION: 123
Azure AD(Azure Active Directory) □□□□ □□□□.
Azure Active Directory □□ □□□□ □□ □□□ □□□□ □□ □□□□ □□□ □□□□□.
□□ □□□ □□□□ □□□ □□□□ □□□□□ □□□.
□□□ □□ □□□ □□□ □□□□ □□□?
A. □ □□□□□ □□□ □□ □□□ □□ □□
B. □ □□□□ □□□ □□ □□□
C. □ □□□□ □□ □□□
D. □ □□□□ □□ □□ □ □□ □□□
E. □ □□□□ □□ □□□ □□□ □□ □□□
Answer: (SHOW ANSWER)
To perform a bulk delete of users in Azure Active Directory, you need to create and upload a CSV file that contains the list of users to be deleted. The file should include the user principal name (UPN) of each user only.
Therefore, the answer is B. The user principal name of each user only. When you use the bulk delete feature in the Azure Active Directory admin center, you need to specify the UPN for each user that you want to delete. The UPN is a unique identifier for each user in Azure AD and is the primary way that Azure AD identifies and manages user accounts. Including additional attributes like the display name or usage location is not required for the bulk delete operation, as the UPN is the only mandatory attribute for the user account.
However, you may include additional attributes in the CSV file if you want to keep track of the metadata associated with each user account.

NEW QUESTION: 124
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.
10□□ □□ □□□□□ □□□□ Azure □□□ □□□□. □□ □□□□□ □□□ □□□ □□□□ □□□□□□.
□□ □□□□ □□□ □□ NSG(□□□□ □□ □□)□ □□ □□□□□.
NSG□ □□□□ □□ □□□□ □□ TCP □□ 8080□ □□□□ □□□□□ □□□□ □□□.
□□ □□: □□□ □□□ □□ □□ □□□ □□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
No, this does not meet the goal. Creating a resource lock and assigning it to the subscription is not enough to ensure that when an NSG is created, it automatically blocks TCP port 8080 between the virtual networks. This is because a resource lock does not affect the configuration or functionality of a resource, but only prevents it from being deleted or modified1. A resource lock does not apply any security rules to an NSG or a virtual network.
To meet the goal, you need to create a custom policy definition that enforces a default security rule for NSGs. A policy definition is a set of rules and actions that Azure performs when evaluating your resources2.

You can use a policy definition to specify the required properties and values for NSGs, such as the direction, protocol, source, destination, and port of the security rule. You can then assign the policy definition to the subscription scope, so that it applies to all the resource groups and virtual networks in the subscription.

NEW QUESTION: 125
VM1□□□ □□ □□□ □□□ Azure □□□ □□□□.
VM1□ CPU □□□□ 80%□ □□□□ □□□ □□□□□ Azure Monitor □□ □□□ □□□□□ □□□.
□□ □□□ User1□ User2□□ □ □□□□□ □□□ □□□□ □□□□ □□□□ □□□.
Azure Monitor□ □□ □□□ □□□□ □□□?
A. □□ □□
B. □□ □□ □□ □□ □□
C. □□ □□
D. Microsoft 365 □□
Answer: A (LEAVE A REPLY)
In Azure Monitor, an alert rule defines the condition (metric or log query) that triggers an alert, but notification and automation actions are managed through an Action Group.
An Action Group is a collection of notification preferences and actions used by Azure Monitor and Service Health alerts. It supports multiple actions such as:
* Email notifications
* SMS messages
* Push notifications
* Webhook triggers
* Automation runbooks
In this scenario, you must send an email message to two users (User1 and User2) whenever the CPU usage of VM1 exceeds 80%. To achieve this, you:
* Create an Action Group in Azure Monitor.
* Add two email receivers (User1 and User2) to that action group.
* Associate the Action Group with the alert rule monitoring VM1's CPU percentage.
Microsoft Official Documentation Extract (Paraphrased):
"Action groups are reusable notification collections that define how you want to be alerted when an alert rule triggers. You can specify multiple email recipients, SMS numbers, and webhook endpoints within one Action Group." (Source: Microsoft Learn - Create and manage action groups in Azure Monitor.)

NEW QUESTION: 126
□□ □□ □□□ □□ □□ □□ □□□ □□□ Azure □□□ □□□□.

Name	SKU
LB1	Basic
LB2	Standard

□□ □□ 6□□ □□□□ □ □□ □□□ □□ □□□ □□□ □□□ □□□□□. □ □□ □□ □□□ □□ □□ 3□□ □□□ □□□□□.
□□□ □□□□ □□ □□ □□□ □□□□ □□□.
□□ □□□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

The virtual machines that will be load balanced by using LB1 must:

- be created in the same availability set or virtual machine scale set.
- be connected to the same virtual network.
- be created in the same resource group.
- be created in the same availability set or virtual machine scale set.
- run the same operating system.

The virtual machines that will be load balanced by using LB2 must:

- be connected to the same virtual network.
- be connected to the same virtual network.
- be created in the same resource group.
- be created in the same availability set or virtual machine scale set.
- run the same operating system.

Answer:

Answer Area

The virtual machines that will be load balanced by using LB1 must:

- be created in the same availability set or virtual machine scale set.
- be connected to the same virtual network.
- be created in the same resource group.
- be created in the same availability set or virtual machine scale set.
- run the same operating system.

The virtual machines that will be load balanced by using LB2 must:

- be connected to the same virtual network.
- be connected to the same virtual network.
- be created in the same resource group.
- be created in the same availability set or virtual machine scale set.
- run the same operating system.

Explanation:

Azure Load Balancers are offered in two SKUs - Basic and Standard - each with distinct capabilities, scalability, and configuration requirements as documented in Microsoft Azure Administrator documentation.

* LB1 - Basic Load Balancer (Basic SKU)The Basic Load Balancer is designed for small-scale, non- production workloads. It has certain limitations compared to the Standard SKU:
* The backend pool of a Basic Load Balancer can only include virtual machines in a single availability set or a single virtual machine scale set.
* This restriction ensures that the Basic Load Balancer maintains consistency and proper failover across VMs sharing the same availability set or scale set.

* It cannot span multiple availability sets or virtual networks.
Therefore, to load balance VMs using a Basic Load Balancer, those VMs must be created in the same availability set or virtual machine scale set.
Correct Option for LB1: be created in the same availability set or virtual machine scale set.
* LB2 - Standard Load Balancer (Standard SKU)The Standard Load Balancer is designed for production- grade workloads and supports advanced features such as zone redundancy, secure by default configuration, and high scalability.
* It supports backend pools composed of virtual machines from different availability zones (in the same region) and multiple virtual machine scale sets.
* The only requirement is that all backend virtual machines must be connected to the same virtual network.
* Standard Load Balancer also supports both public and internal load balancing and provides richer metrics and diagnostic capabilities.
Therefore, for LB2 (Standard), the virtual machines only need to reside within the same virtual network, regardless of their availability sets or zones.
Correct Option for LB2: be connected to the same virtual network.
Microsoft Azure Official Extract (summarized):
"For Basic Load Balancer, backend pool VMs must belong to the same availability set or virtual machine scale set. For Standard Load Balancer, backend pool members can be across availability zones and scale sets, but must be connected to the same virtual network." (From Microsoft Azure Administrator Study Guide - Load Balancer Configuration and Comparison; Azure Docs: azure.microsoft.com > Load Balancer SKUs Comparison.) Final Verified Answer:
LB1: be created in the same availability set or virtual machine scale set
LB2: be connected to the same virtual network

NEW QUESTION: 127
□□ □□ □□□ □□ □□□□□ □□□□ Sub1□ Sub2□□ □ □□ Azure □□□ □□□□.

Name	Location	Subscription	Contains virtual machine
VNet1	East US	Sub1	VM1
VNet2	West US	Sub2	VM2

VM1□ VM2□ □□□ □ □□□ □□ □□□. □□□□ □□□ □□ □□□ □□□□□ □□□. □□□ □□□□ □□□?
A. Azure VPN □□□□□
B. Azure □□ □□
C. □□□ □□ □□(UDR)
D. □□□□ □□□
E. □□□□ □□ □□□□□□(NVA)
Answer: D (LEAVE A REPLY)
Azure Virtual Network (VNet) Peering is the Microsoft-recommended solution to enable seamless communication between two virtual networks, either within the same subscription or across different subscriptions, with minimal cost and administrative overhead.
According to the Microsoft Azure Administrator documentation, VNet Peering connects two virtual networks and allows resources in either VNet to communicate with each other using private IP addresses, just as if they were part of the same network. Traffic between peered VNets remains on the Microsoft backbone network, ensuring low latency, high bandwidth, and no data exposure to the public internet.
There are two types of peering supported:
* VNet Peering (intra-region): For virtual networks in the same Azure region.
* Global VNet Peering: For virtual networks in different Azure regions (for example, East US and West US).
In this scenario, VNet1 (East US, Sub1) and VNet2 (West US, Sub2) are in different regions and different subscriptions. Therefore, Global VNet Peering is the appropriate configuration. Peering can be established across subscriptions, provided that the subscriptions are associated with the same Azure Active Directory tenant or proper permissions exist between tenants.
Unlike VPN gateways or network virtual appliances, VNet peering does not require additional infrastructure or incur data transfer gateway costs, making it the lowest-cost and least administratively complex option.
Once configured, communication between VM1 in VNet1 and VM2 in VNet2 occurs over Microsoft's private backbone network, without public IPs or tunneling.
NEW QUESTION: 128

RG26□□□ □□□ □□□ □□□ Azure □□□ □□□□.

RG26□ □□□ □□□ □□□ □□□□□ □□ □□ □□□□ □□□ □ □□□□□. RG26□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	Type	Location
VM1	Virtual machine	North Europe
RGV1	Recovery Services vault	North Europe
SQLDB01	Azure SQL database	North Europe
AZSQL01	Azure SQL database server	North Europe
sa001	Storage account	West Europe

SQLD01□ RGV1□ □□□□□.

□□□□□ □□□□ Azure Portal□□ RG26□ □□□□□ □□□□□. □□□ □□□□□.

RG26□ □□□□ □□□.

□□ □□ □□□ □□ □□□?

A. SQLDB01 □□□ □□□□□.

B. sa001□ □□□□□.

C. VM1□ □□□□□.

D. VM1□ □□□□□.

Answer: (SHOW ANSWER)

You can't delete a vault that contains backup data. So in this case at first you have to delete the backup of 'SQLD01' before you attempt to delete the vault.

Reference:

https://docs.microsoft.com/en-us/azure/backup/backup-azure-delete-vault

NEW QUESTION: 129

□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.

□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.

Adatum□□□ Azure AD(Azure Active Directory) □□□□ Subscription1□□□ Azure □□□ □□□□. Adatum□□ Developers□□ □□□ □□□□ □□□□. □□1□□ Dev□□ □□□ □□□ □□□□ □□□□.

Dev □□□ □□□□ Azure □□ □□ □□ □ □□ □□□ □□□ □□□ □□□□ □□□.

□□ □□: □□1□□□ □□ □ □□□ □□□ □□□ □□□ □□□□□.

□□□ □□□ □□□□□?

A. □

B. □□□

Answer: (SHOW ANSWER)

The Logic App Operator role only grants the ability to read, enable, disable, and run logic apps. It does not grant the ability to create logic apps. To create logic apps, you need to assign the Logic App Contributor role or a higher-level role such as Owner or Contributor. Then, References: [Built-in roles for Azure resources]

[Azure Logic Apps permissions and access control]

NEW QUESTION: 130

App1□ Azure□ □□□ □□□□□.

□□□□ □□ □□(NSG)□ □□□□□.

□□□□□ App1□ □□ □□ □□□ □□□□ □□□□ □□□□ □□□.

□□□ □□□□ □□□?

A. □□□□□ □□ 443□□ □□□ □□ □□□ □□□□. NSG□ □□ □□□□ □□□□□.

B. □□□□□ □□ 443□□ □□□□ □□ □□□ □□□□. NSG□ □□ □□□□ □□□□□.
C. □□□□□ □□ 443□□ □□□□ □□ □□□ □□□□. NSG□ □ □□□ □□ □□□□ □□□□□.
D. □□□□□ □□ 443□□ □□□ □□ □□□ □□□□. NSG□ □ □□□ □□ □□□□ □□□□□.
Answer: C (LEAVE A REPLY)
As App1 is public-facing we need an incoming security rule, related to the access of the web servers.
Scenario: You have a public-facing application named App1. App1 is comprised of the following three tiers: a SQL database, a web front end, and a processing middle tier.
Each tier is comprised of five virtual machines. Users access the web front end by using HTTPS only.
Topic 3, Contoso Ltd (Consulting Company)Case study
This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.
To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.
At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.
To start the case study
To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.
Overview
General Overview
Contoso, Ltd. is a consulting company that has a main office in Montreal and branch offices in Seattle and New York.
Environment
Existing Environment
Contoso has an Azure subscription named Sub1 that is linked to an Azure Active Directory (Azure AD) tenant. The network contains an on-premises Active Directory domain that syncs to the Azure AD tenant.
The Azure AD tenant contains the users shown in the following table.

Name	Type	Role
User1	Member	None
User2	Guest	None
User3	Member	None
User4	Member	None

Sub1 contains two resource groups named RG1 and RG2 and the virtual networks shown in the following table.

Name	Subnet	Peered with
VNET1	Subnet1, Subnet2	VNET2
VNET2	Subnet1	VNET1, VNET3
VNET3	Subnet1	VNET2
VNET4	Subnet1	None

User1 manages the resources in RG1. User4 manages the resources in RG2.
Sub1 contains virtual machines that run Windows Server 2019 as shown in the following table

Name	IP address	Location	Connected to
VM1	10.0.1.4	West US	VNET1/Subnet1
VM2	10.0.2.4	West US	VNET1/Subnet2
VM3	172.16.1.4	Central US	VNET2/Subnet1
VM4	192.168.1.4	West US	VNET3/Subnet1
VM5	10.0.22.4	East US	VNET4/Subnet1

No network security groups (NSGs) are associated to the network interfaces or the subnets.

Sub1 contains the storage accounts shown in the following table.

Name	Kind	Location	File share	Identity-based access for file share
storage1	Storage (general purpose v1)	West US	sharea	Azure Active Directory Domain Services (Azure AD DS)
storage2	StorageV2 (general purpose v2)	East US	shareb, sharec	Disabled
storage3	BlobStorage	East US 2	**Not applicable**	**Not applicable**
storage4	FileStorage	Central US	shared	Azure Active Directory Domain Services (Azure AD DS)

Requirements

Planned Changes

Contoso plans to implement the following changes:

Create a blob container named container1 and a file share named share1 that will use the Cool storage tier.

Create a storage account named storage5 and configure storage replication for the Blob service.

Create an NSG named NSG1 that will have the custom inbound security rules shown in the following table.

Priority	Port	Protocol	Source	Destination	Action
500	3389	TCP	10.0.2.0/24	Any	Deny
1000	Any	ICMP	Any	VirtualNetwork	Allow

Associate NSG1 to the network interface of VM1.

Create an NSG named NSG2 that will have the custom outbound security rules shown in the following table.

Priority	Port	Protocol	Source	Destination	Action
200	3389	TCP	10.0.0.0/16	VirtualNetwork	Deny
400	Any	ICMP	10.0.2.0/24	10.0.1.0/24	Allow

Associate NSG2 to VNET1/Subnet2.

Technical Requirements

Contoso must meet the following technical requirements:

Create container1 and share1.

Use the principle of least privilege.

Create an Azure AD security group named Group4.

Back up the Azure file shares and virtual machines by using Azure Backup.

Trigger an alert if VM1 or VM2 has less than 20 GB of free space on volume C.
Enable User1 to create Azure policy definitions and User2 to assign Azure policies to RG1.
Create an internal Basic Azure Load Balancer named LB1 and connect the load balancer to VNET1/Subnet1 Enable flow logging for IP traffic from VM5 and retain the flow logs for a period of eight months.
Whenever possible, grant Group4 Azure role-based access control (Azure RBAC) read-only permissions to the Azure file shares.

NEW QUESTION: 131
□□□ □□ □□□ □□□□□ WebApp1□ □□□□ □□□.
Vault1□□ □□ □□□□ □□□ □ □□□?
A. Cert1□
B. Cert1 □□ Cert2□
C. Cert1 □□ Cert3□
D. Cert3 □□ Cert4□ □□
E. Cert1, Cert2, Cert3 □□ Cert4
Answer: A (LEAVE A REPLY)
To meet the technical requirement - "Use TLS for WebApp1" - the web app must be configured with a certificate that is compatible with Azure App Service for HTTPS/TLS binding.
According to the Microsoft Azure Administrator documentation on App Service Certificates and Key Vault integration, the following key points determine which certificates can be used:
* Supported Certificate Format:Azure App Service supports importing certificates in PFX (PKCS #12) format, which includes both the public and private keys necessary for TLS/SSL binding.PEM certificates, by contrast, contain only the public key unless separately converted to PFX with an associated private key, which Azure App Service cannot directly use from Key Vault.
* Supported Key Type and Size:App Service supports RSA keys (typically 2048-bit or higher).Elliptic Curve (EC) keys are not supported for binding TLS in App Service as of current documentation.
* Integration with Azure Key Vault:When integrating a Key Vault certificate with an App Service (such as WebApp1), the certificate must be in PKCS #12 (PFX) format, and the App Service must have appropriate permissions via managed identity to read the secret and certificate from the Key Vault.
From the Vault1 data provided in your scenario:
Name
Content type
Key type
Key size
Cert1
PKCS #12
RSA
2048
Cert2
PKCS #12
RSA
4096
Cert3
PEM
RSA
2048
Cert4
PEM
RSA

4096
Analysis:
* Cert1 and Cert2 are PKCS #12 certificates, so both contain the private key required for TLS.
* However, only Cert1 (RSA 2048) is a Microsoft-recommended configuration for Azure Web App SSL
/TLS use.
* Cert2 has a 4096-bit RSA key. Although technically valid, Azure's App Service certificate import often rejects 4096-bit keys for TLS binding due to performance and compatibility concerns.
* Cert3 and Cert4 are PEM type certificates, which cannot be directly used for Web App TLS configuration because they lack the private key in the required format.
Therefore, according to the Azure Administrator Exam Study Guide and Microsoft official documentation, the only valid certificate that meets the requirements is:
Cert1 only
Final Verified Answer: # A. Cert1 only

NEW QUESTION: 132
Subscription1□□□ Azure □□□ □□□□.
contosostorage□□ Azure Storage □□□ □□ □□ data□□ □□ □□□ □□□□.
□□□ □□ □□□ □□□ □□□□ □□□□□ □□ UNC □□□ □□□□ □□□? □□ □□□□ □□□ □□ □□□ □□□ □□□ □□□□. □ □□ □ □, □□ □ □□ □□ □□□□ □□ □ □□□□. □□□ □□□ □ □□□ □□ □□□ □□□ □□□ □□□□□ □ □□ □□□□.
□□: □□ □□□ 1□□□□.

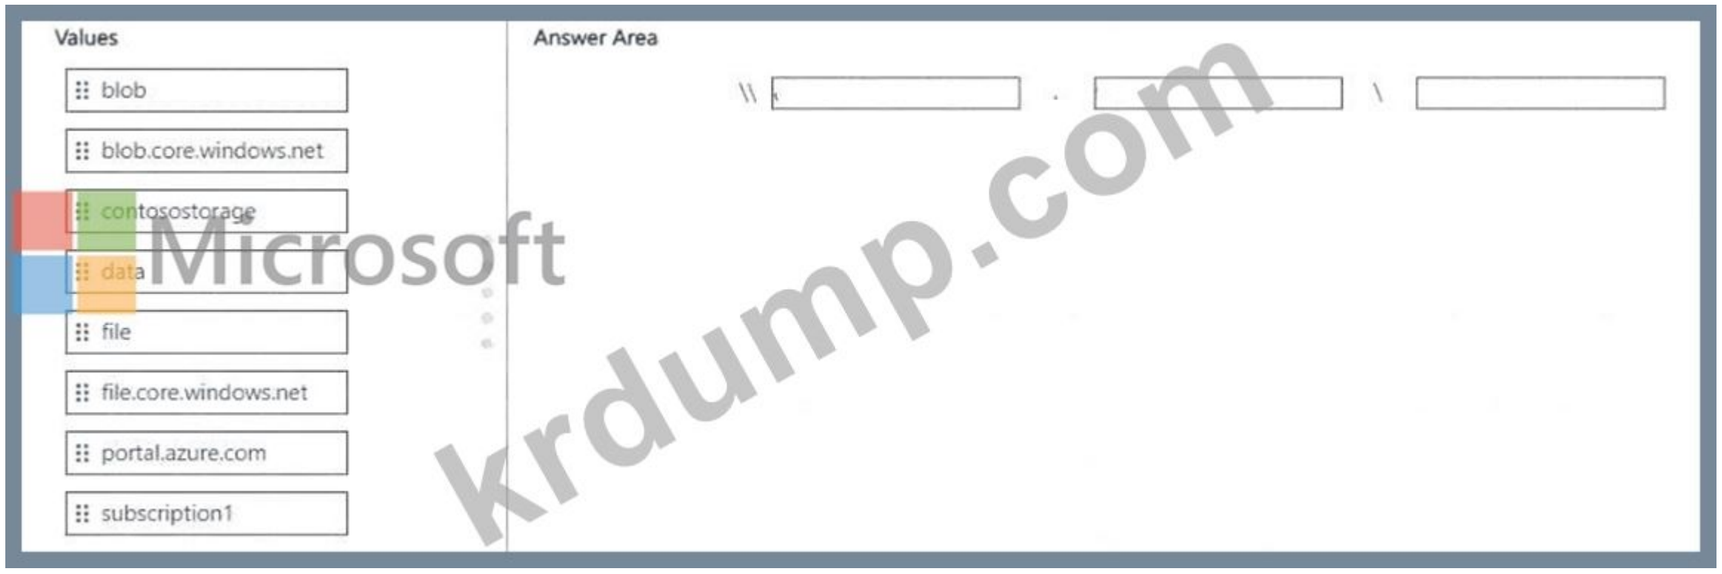

Answer:

Explanation:

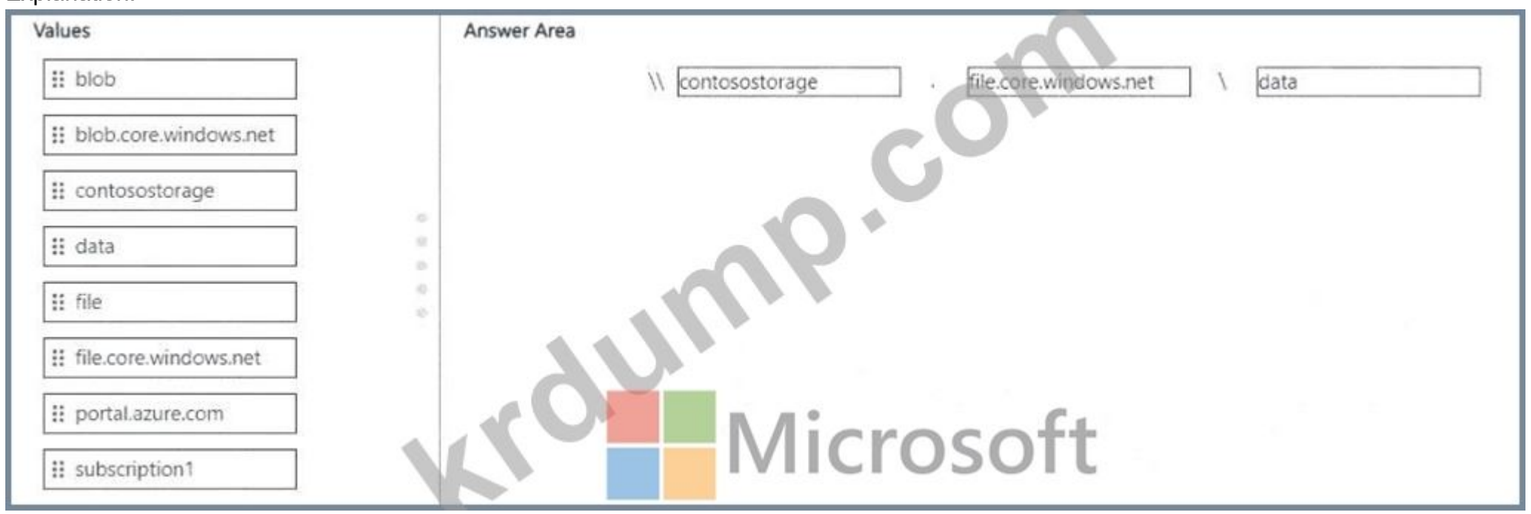

Azure File Shares are accessed using a UNC (Universal Naming Convention) path, which follows a strict and well-defined format in Microsoft Azure. When you create an Azure Storage account and then create a file share within that account, the file share is exposed over the SMB protocol and can be mounted or referenced just like a traditional Windows file share.

According to Microsoft Azure Administrator documentation and study guides, the UNC path format for an Azure file share is:

\\<storage-account-name>.file.core.windows.net\<file-share-name>

In this scenario:

* The Azure subscription name (Subscription1) is not part of the UNC path.
* The storage account name is contosostorage.
* The file share name is data.
* Azure Files always uses the file.core.windows.net endpoint for SMB-based file shares.
* Blob endpoints (blob.core.windows.net) are used only for Blob Storage and are not valid for file shares.

Putting these components together results in the correct UNC path that can be used inside scripts, PowerShell, batch files, or application configurations to reference files stored in the Azure file share.

NEW QUESTION: 133

10□□ Azure □□□ □□□ Azure AD □□□□ □□□□.
□□ □□□□ □□□ □□□ □□□□ □□□□□□ □□□.
□□□ □□□□ □□□?

A. □□ □□ □□

B. Log Analytics □□ □□

C. □□□ □□

D. Azure Application Insights □□□□□

Answer: B (LEAVE A REPLY)

https://learn.microsoft.com/en-us/azure/azure-monitor/essentials/activity-log?tabs=powershell#send-to-log-analytics-workspace Send the activity log to a Log Analytics workspace to enable the Azure Monitor Logs feature, where you: - Consolidate log entries from multiple Azure subscriptions and tenants into one location for analysis together.

NEW QUESTION: 134

□□ □□□□ □□□□ Contoso.com□□□ Azure Active Directory □□□□ □□□□.

Name	Role
User1	Cloud device administrator
User2	User administrator

Contoso.com□□ □□ Windows 10 □□□ □□□□□.

Name	Join type
Device1	Azure AD registered
Device2	Azure AD joined

Contoso.com□□ □□ □□ □□□ □□□□.

Name	Join type	Owner
Group1	Assigned	User1
Group2	Dynamic Device	User2

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

Box 1: Yes

User1 is a Cloud Device Administrator.

Device2 is Azure AD joined.

Group1 has the assigned to join type. User1 is the owner of Group1.

Note: Assigned groups - Manually add users or devices into a static group.

Azure AD joined or hybrid Azure AD joined devices utilize an organizational account in Azure AD Box 2: No User2 is a User Administrator.

Device1 is Azure AD registered.

Group1 has the assigned join type, and the owner is User1.

Note: Azure AD registered devices utilize an account managed by the end user, this account is either a Microsoft account or another locally managed credential.

Box 3: Yes

User2 is a User Administrator.

Device2 is Azure AD joined.

Group2 has the Dynamic Device join type, and the owner is User2.

References:

https://docs.microsoft.com/en-us/azure/active-directory/devices/overview

NEW QUESTION: 135

VM1□ VM4□ □□□ □ □□□ □□□□ □□□. □□□□ □□ □□□ □□□□□ □□□.

□□□ □□ □□□?

A. VNET1□□ VNET3□ □□ □□□ □□ □□□ □□□□.

B. VM4□ IP □□ 10.0.1.5/24□ □□□□□.

C. VNET1□ VNET3 □□ □□□□ □□□□□.

D. NSG□ □□□□ NSG□ VMI □ VM4□ □□□□□.

Answer: (SHOW ANSWER)

To enable communication between virtual machines (VMs) located in different virtual networks (VNets) in Azure, the most efficient and recommended approach-according to Microsoft Azure Administrator documentation-is to use VNet peering.

1. Background and Scenario Analysis

From the case study:

VM1 and VM4 are located in different VNets (VNET1 and VNET3).

The requirement is to ensure that VM1 can communicate with VM4.

The solution must minimize administrative effort and cost.

2. Microsoft Documentation Insight: VNet Peering

According to Microsoft Learn: "Virtual network peering":

"Virtual network peering seamlessly connects two Azure virtual networks. The virtual networks appear as one for connectivity purposes. Traffic between peered virtual networks uses private IP addresses, as if they were part of the same network, and the traffic stays entirely on the Microsoft backbone network." Key characteristics of VNet peering:

Enables private IP connectivity between resources across peered VNets.

No need to deploy or maintain gateways (unlike VPN gateways).

Provides low latency and high bandwidth.

Supports transitive routing through additional configurations.

Minimal administrative overhead - peering can be created with just a few clicks or PowerShell/CLI commands.

3. Why the Other Options Are Incorrect

A). Create a user-defined route (UDR) from VNET1 to VNET3.

A UDR alone cannot enable connectivity between VNets unless a gateway or peering already exists.

Without a connection path, a route has no effect.

B). Assign VM4 an IP address of 10.0.1.5/24.

This would attempt to place VM4 in the same subnet as VM1, but cross-VNet subnet IP assignment is not possible in Azure. Each VNet has its own isolated address space.

D). Create an NSG and associate it with VM1 and VM4.

Network Security Groups control traffic filtering within or between existing network connections. They do not create connectivity between isolated VNets.

4. Why Peering Is the Correct and Simplest Solution

Establishing VNet peering between VNET1 and VNET3 will:

Instantly enable bidirectional private IP communication between VM1 and VM4.

Minimize administrative effort (no gateways, routing tables, or IP reconfiguration).

Maintain security and performance through Microsoft's internal backbone.

Avoid additional costs compared to deploying VPN gateways.

5. Implementation Summary

Steps to configure:

In the Azure Portal, go to VNET1 # Peerings # Add.

Choose VNET3 as the peer virtual network.

Enable Allow virtual network access in both directions.

Once completed, both VMs (VM1 and VM4) will communicate using their private IPs.

Final Verified Answer: # C. Establish peering between VNET1 and VNET3

References (Microsoft Official Documentation):

Microsoft Learn - Virtual network peering overview

Microsoft Learn - Create, change, or delete a virtual network peering

Microsoft Learn - Azure virtual network connectivity options and recommendations

NEW QUESTION: 136
Subscription1□□□ Azure □□□ □□□□. Subscription1□□ □□ □□ □□ □□ □□□ □□□ □□□□ □□□□.

Name	Azure region	Assigned Azure Policy
RG1	West Europe	Policy1
RG2	North Europe	Policy2
RG3	France Central	Policy3

RG1□□ WebApp1□□□ □ □□ □□□□. WebApp1□ □□□□ □□□□.
WebApp1□ RG2□ □□□□□.
□ □□□ □□□ □□□□?
A. WebApp1□ □□ App Service □□□ □□□□ □□□□□. Policy2□ WebApp1□ □□□□□.
B. WebApp1□ □□ App Service □□□ □□□□ □□□□□. Policy1□ WebApp1□ □□□□□.
C. WebApp1□ App Service □□□ □□□□□ □□□□□. □□ 2□ WebApp1□ □□□□□.
D. WebApp1□ App Service □□□ □□□□□ □□□□□. Policy1□ WebApp1□ □□□□□.
Answer: C (LEAVE A REPLY)

NEW QUESTION: 137
□□ □□ □□□ □□ □□□ □□□ Azure □□□ □□□□.
□□□□□:void(0)

Name	Public IP SKU	Connected to	Status
VM1	*None*	VNET1/Subnet1	Stopped (deallocated)
VM2	Basic	VNET1/Subnet2	Running

□□ □□□ □□ □□ □□□□ □□□□□.
* □□: LB1
* □□□
* SKU: □□
* □□ □□□□ VNET1
LB1□ □□□ □□ VM1□ VM2□ □□□ □ □□□ □□□□ □□□.
□□ □□: □□ SKU □□ IP □□□ □□□□, □□ □□□ VM1□ □□□□ □□□□□□ □□□ □□ VM1□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
You can only attach virtual machines that are in the same location and on the same virtual network as the LB.

Virtual machines must have a standard SKU public IP or no public IP.
The LB needs to be a standard SKU to accept individual VMs outside an availability set or vmss. VMs do not need to have public IPs but if they do have them they have to be standard SKU. Vms can only be from a single network. When they don't have a public IP they are assigned an ephemeral IP.
Also, when adding them to a backend pool, it doesn't matter in which status are the VMs.
Note: Load balancer and the public IP address SKU must match when you use them with public IP addresses.

NEW QUESTION: 138
□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type	Description
vm1	Virtual machine	Uses a basic public IP address
vm2	Virtual machine	Uses a basic public IP address
nsg1	Network security group (NSG)	Allows incoming traffic to port 443
lb1	Azure Standard Load Balancer	*None*

lb1□ □□□□ vm1 □ vm2□ □□ HTTPS □□□ □□□ □□□□ □□□.
□□ □ □□ □□□ □□□□ □□□□ □□□? □□□□ □□ □□□□ □□ □□□ □□ □□□□ □□ □□□ □□□□ □□□□□.

Answer:

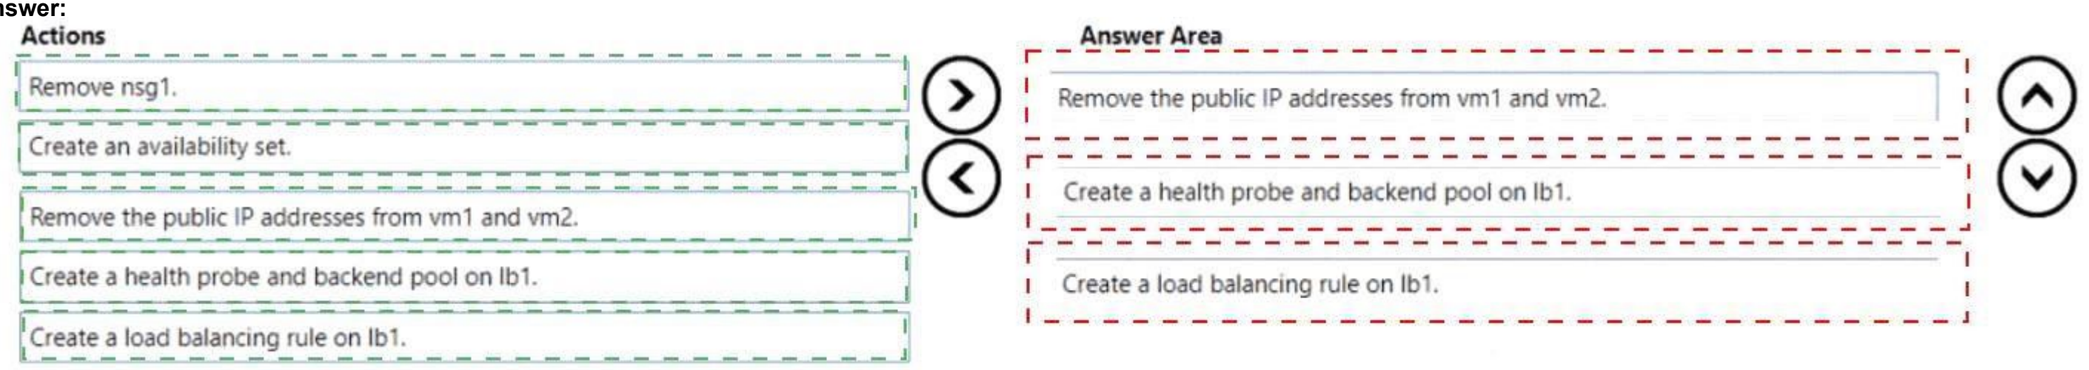

Explanation:

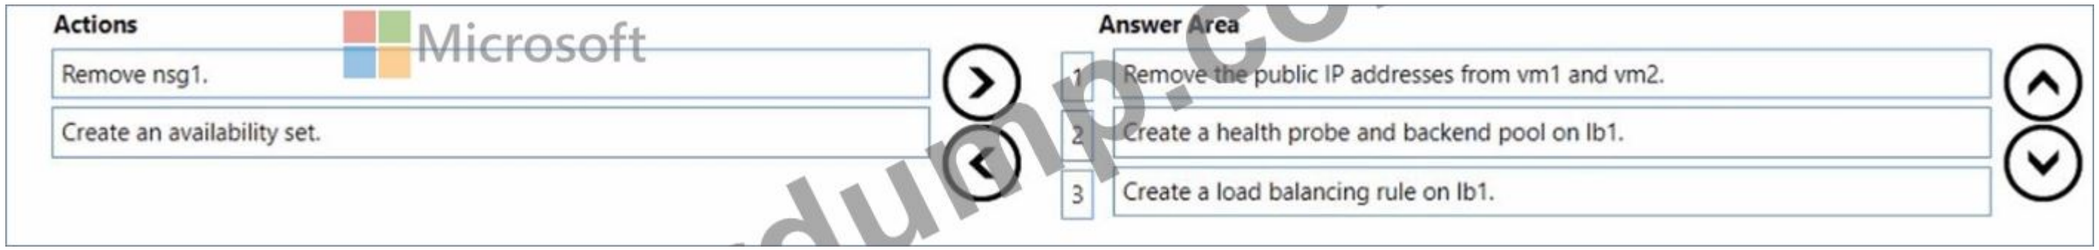

To configure an Azure Standard Load Balancer to distribute HTTPS (TCP port 443) traffic between two virtual machines (vm1 and vm2), several specific configuration steps must be followed according to Microsoft Azure Administrator documentation ("Load Balancer Standard SKU - configuration and differences").

* Remove the public IP addresses from vm1 and vm2:Azure Standard Load Balancer only supports backend resources that use private IP addresses. Virtual machines that are part of a Standard Load Balancer backend pool must not have their own Basic Public IP addresses. Therefore, both vm1 and vm2 must have their public IPs removed before they can be added to the load balancer backend pool.
* Create a health probe and backend pool on lb1:A health probe is required to monitor the availability of backend VMs. The backend pool defines which VMs receive load-balanced traffic. In this case, vm1 and vm2 are added to the backend pool.
* Create a load balancing rule on lb1:The load balancing rule defines how traffic is distributed - specifying frontend IP configuration, protocol (TCP), port (443), backend pool, and health probe.

This configuration allows lb1 to evenly distribute HTTPS traffic across vm1 and vm2, ensuring high availability and secure connectivity using the Standard Load Balancer's internal/private endpoints.

NEW QUESTION: 139

VM1□□□ □□ □□□ □□□ Azure □□□ □□□□.

DC1□□□ □□□ □□□□□ □□□ □□□□ □□□ □□□ □□□□. ExpressRoute□ □□□□ □□□ □□□ Azure□ □□□□ □ □□□□□.

VM1□ DC1 □□ □□□□ □□ □□□ □□□□□ Connection Monitor□ □□□□ □□□.

DC1□ □□□ □□□□ □□□?

A. Log Analytics □□□□

B. Azure Network Watcher Agent □□ □□ □□

C. Azure Monitor □□□□ □□

D. Azure Arc □□ □□□ Azure Connected Machine □□□□

Answer: D (LEAVE A REPLY)

This question focuses on how to monitor network connectivity and latency between Azure virtual machines and on-premises resources using Azure Network Watcher - Connection Monitor (v2).

Scenario Breakdown

VM1: Azure virtual machine (in your subscription)

DC1: On-premises domain controller (in your datacenter)

Connectivity: Via ExpressRoute

Goal: Use Connection Monitor to track network latency between VM1 (Azure) and DC1 (on-premises) To achieve this, both endpoints (VM1 and DC1) must have agents capable of collecting and sending network telemetry data to Azure Monitor.

Understanding Azure Connection Monitor (v2)

According to Microsoft Learn ("Monitor network connectivity with Connection Monitor"):

"Connection Monitor (v2) enables you to monitor network connectivity between Azure and on-premises resources. You can monitor connections between Azure VMs, Azure Arc-enabled servers, and any endpoint reachable over TCP or ICMP." To participate in a hybrid connection, the on-premises machine must be onboarded to Azure Arc.

Azure Arc connects your non-Azure servers to Azure Resource Manager and allows management and monitoring using Azure services, including Network Watcher's Connection Monitor.

Required Agent for On-Premises Monitoring

For on-premises servers (like DC1), Azure Arc uses the Azure Connected Machine agent (formerly known as the Azure Arc agent).

This agent:

Registers the on-premises machine as an Azure Arc-enabled server in Azure.
Enables the use of Azure Monitor, Defender for Cloud, Update Management, and Connection Monitor on that server.
Once the Azure Connected Machine agent is installed and the server is connected through Azure Arc, you can add it as a source or destination endpoint in Connection Monitor to measure latency and packet loss.
Why Other Options Are Incorrect
Option
Description
Why Incorrect
A). Log Analytics agent
Used for data collection (logs and metrics) for Azure Monitor.
Does not support Connection Monitor (v2) endpoint monitoring. Deprecated in favor of Azure Monitor Agent.
B). Azure Network Watcher Agent extension
Used only on Azure VMs, not on on-premises servers.
Cannot be installed on DC1 (non-Azure).
C). Azure Monitor agent extension
Used for telemetry/log ingestion into Azure Monitor.
Does not support connectivity monitoring for hybrid endpoints.
D. Azure Connected Machine agent
Connects on-premises servers to Azure Arc, enabling Connection Monitor and other Azure services.
Correct and required.
Verification (Microsoft Documentation Extract)
From Microsoft Learn - "Monitor hybrid connectivity using Connection Monitor":
"To monitor on-premises resources, onboard the servers to Azure Arc and install the Azure Connected Machine agent. This agent enables Connection Monitor to collect network connectivity data from on-premises endpoints."
Final Verified Answer:
D). the Azure Connected Machine agent for Azure Arc-enabled servers
Summary of Key Points
Connection Monitor (v2) can track network latency between Azure and on-premises systems.
On-premises servers must be Azure Arc-enabled using the Azure Connected Machine agent.
Azure VMs use the Network Watcher extension, while non-Azure machines require Azure Arc for integration.
Correct Answer: D. the Azure Connected Machine agent for Azure Arc-enabled servers

NEW QUESTION: 140
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.
VM1 □ VM2□□ □ □□ Azure □□ □□□ □□□ App1□□□ □□ □□□□.
App1□ □□ □□□ Azure Load Balancer□ □□□□ □□□□□.
VM2□ □□□□ □□□□ □□ □□□ □□ □□□ □□ □□□□.

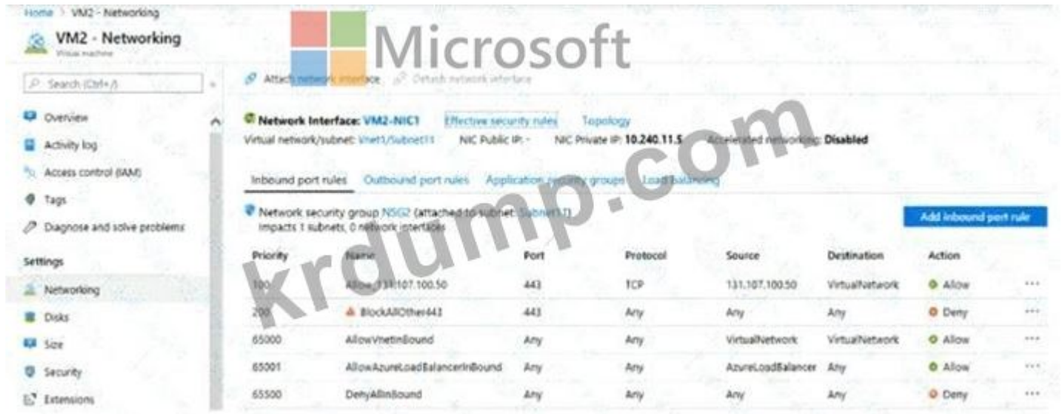

TCP □□ 443□ □□ 131.107.100.50□□ App1□ □□ □□□ □□□□□ □□□□□□. Load Balancer □□□ □□□□ □□□□□□ □□□□□.
TCP □□ 443□ □□ 131.107.100.50□□ App1□ □□ □□□ □□□□□ □□□ □ □□□ □□□□ □□□.
□□ □□: 131.107.100.50 □□□ □□ □□□□ □□□□ □□□ 64999□ □□□□ □□ □□□ □□□□□.
□□□ □□□ □□□□□?

A. □

B. □□□

Answer: (SHOW ANSWER)

In the provided NSG configuration for VM2, the effective inbound security rules show the following priorities:

Priority

Name

Port

Protocol

Source

Action

100

Allow_131.107.100.50

443

TCP

131.107.100.50

Allow

200

BlockAllOther443

443

TCP

Any

Deny

65000+

Default rules

Various

Any

Allow/Deny
The problem states that connections from 131.107.100.50 to port 443 fail, even though the Load Balancer rules are correct.
The proposed solution is to create a new inbound rule that denies all traffic from 131.107.100.50 with a priority of 64999.
However, this does not solve the problem - it makes it worse, because:
* The rule with priority 64999 has a higher number, meaning it is evaluated later than existing rules.
* Since NSG rules are processed in ascending order of priority, the first matching rule applies.
* The rule at priority 200 (BlockAllOther443) already denies the connection before the new rule is even evaluated.
To fix the issue, the correct solution would be to modify or remove the "BlockAllOther443" rule, or adjust priorities so that the allow rule for 131.107.100.50 is evaluated first.
Simply adding another deny rule with a lower priority number (higher numeric value) will not override existing denies.
Final Verified Answer: B. No

NEW QUESTION: 141
□□ □□ 2 □□□ VNET□□□ □□ □□□□□ □□□ Azure □□□ □□□□. VM1-NI□□ □□□□ □□□□□□ VNET1□ □□□□ □□□□.
□□ Azure Resource Manager □□□□ □□□□□ □□□□□□.

```
{
    "apiVersion": "2017-03-30",
    "type": "Microsoft.Compute/virtualMachines",
    "name": "VM1",
    "zones": "1",
    "location": "EastUS2",
    "dependsOn": [
      "[resourceId('Microsoft.Network/networkInterfaces', 'VM1-NI')]"
    ],
    "properties": {
      "hardwareProfile": {
        "vmSize": "Standard_A2_v2"
      },
      "osProfile": {
        "computerName": "VM1",
        "adminUsername": "AzureAdmin",
        "adminPassword": "[parameters('adminPassword')]"
      },
      "storageProfile": {
        "imageReference": "[variables('image')]",
        "osDisk": {
          "createOption": "FromImage"
        }
      },
      "networkProfile": {
        "networkInterfaces": [
          {
            "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM1-NI')]"
          }
        ]
      }
    },
    {
      "apiVersion": "2017-03-30",
      "type": "Microsoft.Compute/virtualMachines",
      "name": "VM2",
      "zones": "2",
      "location": "EastUS2",
      "dependsOn": [
        "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
      ],
      "properties": {
        "hardwareProfile": {
          "vmSize": "Standard_A2_v2"
        },
        "osProfile": {
          "computerName": "VM2",
          "adminUsername": "AzureAdmin",
          "adminPassword": "[parameters('adminPassword')]"
        },
        "storageProfile": {
          "imageReference": "[variables('image')]",
          "osDisk": {
            "createOption": "FromImage"
          }
        },
        "networkProfile": {
          "networkInterfaces": [
            {
              "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
            }
          ]
        }
      }
    }
```

Answer:

Explanation:

"A resource can only be created in a virtual network that exists in the same region and subscription as the resource." https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-vnet-plan-design- arm#regions

NEW QUESTION: 142

□□ □□□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

Answer:

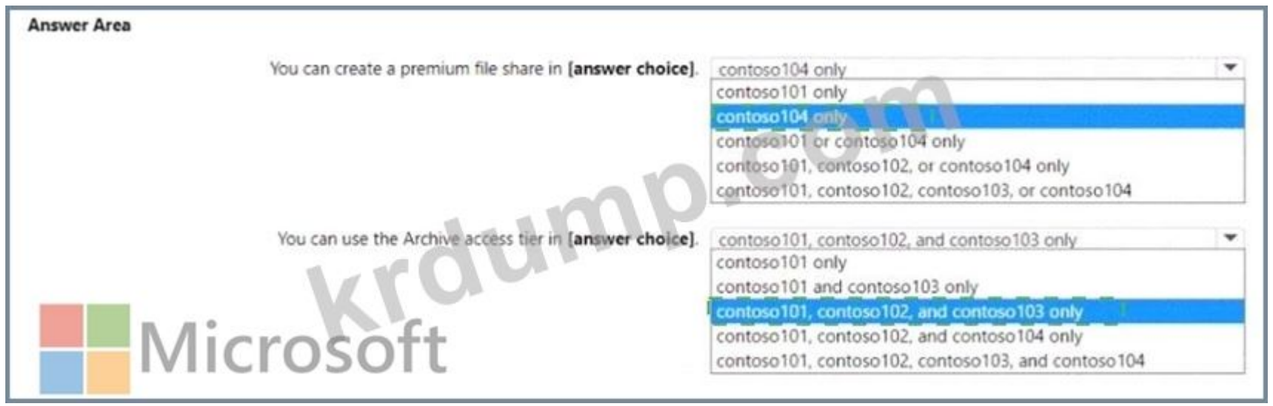

Explanation:

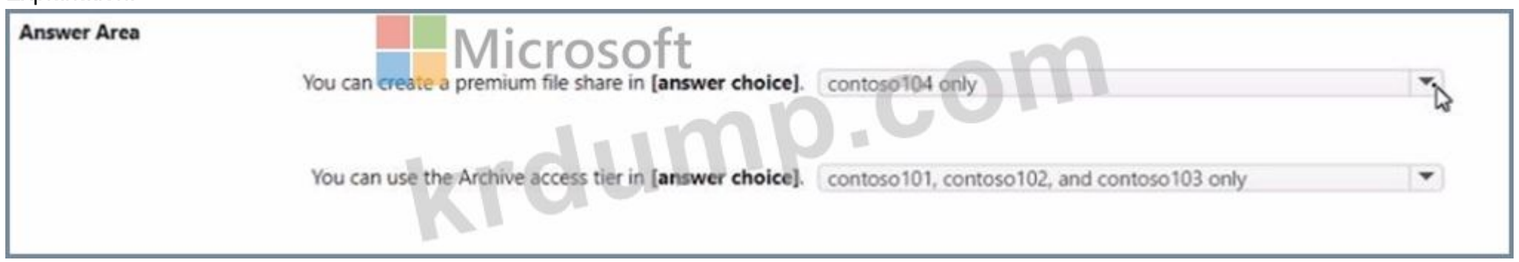

Scenario:

You have the following storage accounts:

Name

Type

Kind

Resource Group

Location

contoso101

Storage account

StorageV2

RG1

East US
contoso102
Storage account
Storage
RG1
East US
contoso103
Storage account
BlobStorage
RG1
East US
contoso104
Storage account
FileStorage
RG1
East US
Question 1:
You can create a Premium file share in ___ ?
Answer: contoso104 only
Premium file shares in Azure use Azure Files Premium performance tier, which is available only with FileStorage accounts.
Per Microsoft Azure documentation:
"To create a Premium file share, you must use a storage account of kind FileStorage with performance tier set to Premium." contoso104 is the only account of type FileStorage, so only it supports premium file shares.
Question 2:
You can use the Archive access tier in ___ ?
Answer: contoso101, contoso102, and contoso103 only
The Archive access tier is available only for Blob storage objects stored in:
General-purpose v2 (StorageV2) accounts, or
BlobStorage accounts.
It is not supported in:
FileStorage accounts (used for file shares).
Thus:
contoso101 (StorageV2) # Supports Archive tier
contoso102 (Storage) # Legacy (no Archive support)
contoso103 (BlobStorage) # Supports Archive tier
contoso104 (FileStorage) # File shares only
Correction: Storage (V1) supports only Hot and Cool access tiers, not Archive.
So, the correct accounts that support the Archive tier are:
contoso101 and contoso103 only
Final Verified Answers:
Question
Correct Answer

You can create a premium file share in
contoso104 only
You can use the Archive access tier in
contoso101 and contoso103 only
Microsoft Azure Administrator Documentation Extracts:
"Premium file shares are available only in FileStorage accounts and provide low-latency, high-performance storage for Azure Files."
"The Archive tier is supported for Blob Storage and General-purpose v2 accounts only." (Source: Microsoft Learn - "Azure storage account overview" & "Access tiers for Azure Blob Storage")

NEW QUESTION: 143
Azure Blob Storage□ Azure File Storage□ □□□□ storage1□□□ Azure Storage □□□ □□□□.
AzCopy□ □□□□ storage1□ Blob □□□□ □□ □□□□ □□□□ □□□□ □□□.
□ □□□ □□□ □□ □□ □□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

Blob storage:

- Azure Active Directory (Azure AD) only
- Shared access signatures (SAS) only
- Access keys and shared access signatures (SAS) only
- Azure Active Directory (Azure AD) and shared access signatures (SAS) only
- Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

File storage:

- Azure Active Directory (Azure AD) only
- Shared access signatures (SAS) only
- Access keys and shared access signatures (SAS) only
- Azure Active Directory (Azure AD) and shared access signatures (SAS) only
- Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

Answer:

Blob storage:

- Azure Active Directory (Azure AD) only
- Shared access signatures (SAS) only
- Access keys and shared access signatures (SAS) only
- Azure Active Directory (Azure AD) and shared access signatures (SAS) only
- Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

File storage:

- Azure Active Directory (Azure AD) only
- Shared access signatures (SAS) only
- Access keys and shared access signatures (SAS) only
- Azure Active Directory (Azure AD) and shared access signatures (SAS) only
- Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

Explanation:

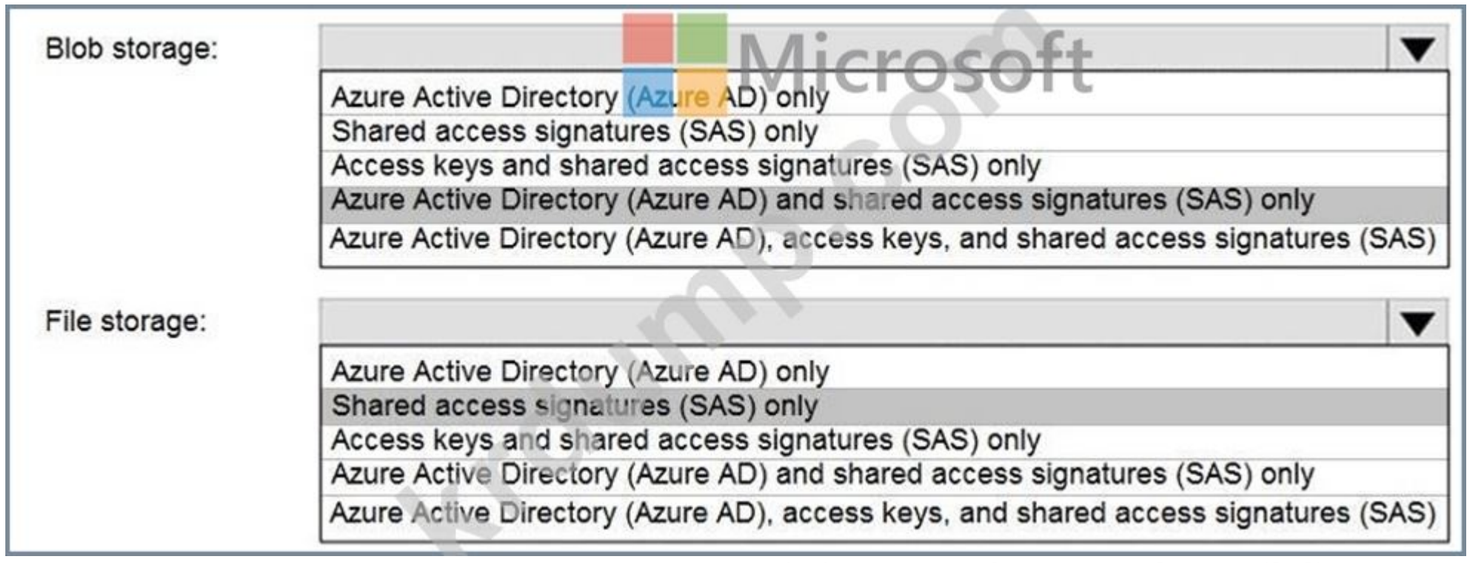

You can provide authorization credentials by using Azure Active Directory (AD), or by using a Shared Access Signature (SAS) token.

Box 1:

Both Azure Active Directory (AD) and Shared Access Signature (SAS) token are supported for Blob storage.

Box 2:

Only Shared Access Signature (SAS) token is supported for File storage.

Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10

NEW QUESTION: 144

Backup1□□□ □□□ Azure Backup □□ □□ □□, Recovery1□□□ □□□ Recovery Services □□ □□ □□ □ □□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

□□□□ □□□ □□□□□.

Backup1□ □□□ □ □□ □□□□ Recovery1□ □□□ □ □□ □□□□ □□□□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

Name	Type
VM1	Virtual machine
Disk1	Disk
App1	Azure App Service web app
DB1	Azure SQL Database

□□: □□ □□□ 1□□□□.

Answer:

Explanation:

Azure provides two different vault types for backup, each designed for specific workloads: Azure Backup vaults and Recovery Services vaults. Selecting the correct vault depends on the resource type being protected.

An Azure Backup vault is the newer vault type introduced to support Azure Disk Backup. According to the Azure Administrator documentation, Backup vaults are specifically designed for Azure managed disks and use a simplified policy model. They do not support backing up full virtual machines, Azure App Service apps, or Azure SQL Database (PaaS). Therefore, from the listed resources, Disk1 (Disk) is the only supported resource that can be backed up to Backup1.

A Recovery Services vault is the traditional and more feature-rich vault used to back up Azure Virtual Machines, as well as workloads such as SQL Server running inside Azure VMs. The documentation clearly states that Azure VMs must be backed up using a Recovery Services vault, not a Backup vault. As a result, VM1 (Virtual machine) must be backed up to Recovery1.

The remaining resources are not supported in either vault for this scenario:

* App1 (Azure App Service web app) uses App Service backup, not Azure Backup.
* DB1 (Azure SQL Database) uses built-in Azure SQL backup, not Recovery Services or Backup vaults.

Final Verified Answers:

* Backup1 # Disk1
* Recovery1 # VM1

NEW QUESTION: 145

□□□□□ Microsoft Entra □□□□ □□□ Azure □□□ □□□□. □□ □□□□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	On-premises sync enabled
User1	No
User2	Yes

□□ □□ □□□ Azure Files □□□ □□□□.

Name	Storage account
share1	contoso2024
share2	contoso2024
share3	contoso2025

□□ □□□ □□ contoso2024□ □□ ID □□ □□□□ □□□□□.

□□ □ □□□ □□, □□□ □□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□. □□: □□□ 1□□□□.

Answer Area

Statements	Yes	No
User1 can access the content in share1.	○	○
User2 can access the content in share2.	○	○
User2 can access the content in share3.	○	○

Answer:

Answer Area

Statements	Yes	No
User1 can access the content in share1.	○	○
User2 can access the content in share2.	○	○
User2 can access the content in share3.	○	○

Explanation:

This question examines your understanding of Azure Files identity-based authentication and Active Directory integration for Azure file shares.

Let's analyze it step by step using official Microsoft Azure Administrator (AZ-104) documentation concepts.

1. Azure Files Identity-Based Access Overview

Azure Files supports identity-based authentication and authorization through:

On-premises Active Directory Domain Services (AD DS)

Azure Active Directory Domain Services (Azure AD DS)

Azure AD Kerberos (hybrid identities required)

Important Note (Microsoft Docs):

"Azure Active Directory (Azure AD) is not a domain controller. Azure AD-only accounts are not supported for SMB access to Azure file shares." This means Azure AD cloud-only users (not hybrid) cannot access SMB file shares using identity-based access.

2. Identity Sync (Hybrid Setup)

User

On-premises Sync Enabled

User1

No

User2

Yes

User1 is a cloud-only user (not hybrid).

Cannot authenticate using SMB to an Azure Files share because only users synchronized from on-premises AD (hybrid users) are supported.

User2 is synchronized from on-premises AD (hybrid).

Can authenticate using SMB and access identity-based file shares integrated with AD DS.

3. Storage Account Configuration

Share

Storage Account

share1

contoso2024

share2
contoso2024
share3
contoso2025
contoso2024
Configured with Active Directory (AD DS) integration (see exhibit).
Default share-level permissions are enabled for all authenticated users and groups with the Storage File Data SMB Share Contributor role.
This means any authenticated domain user (hybrid) has access.
contoso2025
No indication of AD DS configuration in the scenario.
Hence, it is not configured for identity-based access.
4. Step-by-Step Validation
User1 and share1 (contoso2024)
User1 is not hybrid (no on-prem sync).
SMB authentication requires Kerberos via domain-joined identity.
Result: # Cannot access share1.
User2 and share2 (contoso2024)
User2 is hybrid (on-prem sync enabled).
contoso2024 supports AD DS integration and allows authenticated domain users.
Result: # Can access share2.
User2 and share3 (contoso2025)
contoso2025 is not configured for AD DS integration.
Without AD DS/AD DS Kerberos setup, SMB access using identity is not possible.
Result: # Cannot access share3.
Official Microsoft Extract (from Azure Files identity-based authentication guide):
"Azure file shares only support SMB access for users and devices that are authenticated by an Active Directory domain controller.
Azure AD-only users are not supported.
You must have hybrid identities synchronized from Active Directory using Azure AD Connect."
"If identity-based access is enabled, all domain-joined and authenticated users with assigned roles (such as Storage File Data SMB Share Contributor) can access file shares."

NEW QUESTION: 146
VM1□□□ Azure □□ □□□ □□□□. VM1□ ARMljson□□□ □□□ □□ Azure Resource Manager □□□□ □□□□ □□□□□□□.
VM1□ □□ □□□ □□□ □□□□ □□□ □□□□□.
VM1□ □□ □□ □□□□ □□□□ □□□.
□□ □□: VM1 □□□ + □□□ □□□□□□ □□□□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: (SHOW ANSWER)
When Microsoft schedules maintenance that may affect a virtual machine (VM), administrators have two options:
* Wait for Azure to perform the maintenance automatically.

* Move the VM to a new host proactively to minimize or avoid downtime.
According to the Microsoft Azure Administrator documentation, using the "Redeploy + reapply" option in the Azure portal forces the VM to migrate to a new Azure host within the same region and availability set.
Here's what happens when you click Redeploy:
* Azure shuts down the existing VM.
* The VM is moved to a new physical host within the same region.
* The operating system disk and data disks are retained, and network configuration (NICs, public IP, etc.) is preserved.
* Once redeployed, Azure reassigns the VM to the same virtual network and reattaches disks and configurations.
This method is explicitly recommended by Microsoft to proactively address maintenance events or host- related issues like "unresponsive VM" or "performance degradation." Therefore, using the Redeploy operation from the "Redeploy + reapply" blade effectively achieves the goal of moving the VM to a different host immediately.

NEW QUESTION: 147
storage1□□□ Azure Storage □□□ □□□, □ □□□ containers□□ Blob □□□□□ □□□□ □□□□. container!□ □□□ □ □□□□ 1□ □□ □□□□ □□□ □□□□ □□□. □□□ □□□□ □□□?
A. □□□ □□
B. □□□ □□(□□ 1□) □□
C. □□□ □□
D. □□□ □□
Answer: (SHOW ANSWER)
To prevent new content added to an Azure Blob container from being modified or deleted for a specific duration (in this case, one year), you must configure an immutability policy by setting a time-based retention policy or legal hold within an access policy on the container.
According to Microsoft's Azure Storage documentation, immutability policies are configured under
"Immutable blob storage", which allows you to store data in a WORM (Write Once, Read Many) state. When a time-based retention policy is set (for example, one year), any blob data added to that container cannot be modified or deleted until the retention period expires.
The configuration is done by defining an access policy on the container and specifying parameters like retentionPeriodInDays. This ensures compliance with regulations such as SEC 17a-4(f), CFTC 1.31(d), and FINRA Rule 4511.
Other options such as access level, IAM settings, and access tier control visibility, permissions, or storage costs but do not enforce immutability or write protection.
Hence, to achieve the goal of preventing modifications to new blobs for one year, you must configure an access policy with a time-based immutability (retention) policy on the container.
Final Verified Answer: C. an access policy

NEW QUESTION: 148
VM1□□□ Azure □□ □□□ □□□□.
Azure□ VM1□□ □□□□ □□□□□.
VM1□ □□□ □□□ □□□ □□□ □□□□ □□□□□ □□□ □□ Azure Monitor□□ □□ □□□ □□□□.
□□□□□ □□□ □□□ □□□□ □□□.
□□□ □□□□ □□□?
A. □□□ □□
B. Azure Log Analytics □□ □□
C. □□ □□
D. □□ □□ □□
Answer: B (LEAVE A REPLY)
Azure Monitor can collect data directly from your Azure virtual machines into a Log Analytics workspace for analysis of details and correlations. Installing the Log Analytics VM extension for Windows and Linux allows Azure Monitor to collect data from your Azure VMs.

Azure Log Analytics workspace is also used for on-premises computers monitored by System Center Operations Manager.
Reference:
https://docs.microsoft.com/en-us/azure/azure-monitor/learn/quick-collect-azurevm

NEW QUESTION: 149
□□□ □□□□□ □□ □□□ □□ □□□ □□□□ □□□.
□ □□□□ □□ Azure □□□□ □□□ □ □□□? □□ □□□□ □□□ □□□ □□□□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer:

Explanation:

In Microsoft Azure, containerized application deployment can occur using multiple services, depending on the image type and operating system platform used. The question refers to two images - Image1 (Windows Server) and Image2 (Linux) - which are stored in an Azure Container Registry (ACR).

According to the Microsoft Azure Administrator Study Guide (AZ-104) and official Azure documentation, both Windows-based and Linux-based container images can be deployed using any of the following services, depending on the workload requirements:

* Azure App Service (Web App for Containers) - supports both Windows and Linux containers for web applications. It allows developers to directly deploy containerized applications from Docker Hub, Azure Container Registry, or a private registry.
* Azure Container Apps - serverless container hosting designed for microservices and event-driven architectures. It supports Linux and Windows containers, using Kubernetes behind the scenes, without requiring users to manage the cluster infrastructure.
* Azure Container Instances (ACI) - provides lightweight, serverless containers for quick deployment and isolated workloads. It supports both Windows and Linux images and can pull directly from Azure Container Registry.

Therefore, since both Image1 (Windows Server) and Image2 (Linux) are valid container images supported by the same three Azure container services, the correct and verified solution is to select "App Service, Azure Container Apps, or Azure Container Instances" for both.

This matches Azure documentation under:

* "Run containerized applications in Azure App Service"
* "Deploy containers using Azure Container Instances"
* "Build and deploy microservices using Azure Container Apps"

Each of these Azure services supports containerized deployments from ACR regardless of the underlying operating system image type.

Final Verified Answer:

Image1: App Service, Azure Container Apps, or Azure Container Instances

Image2: App Service, Azure Container Apps, or Azure Container Instances

Topic 2, Contoso LtdOverview

Contoso, Ltd. is a manufacturing company that has offices worldwide. Contoso works with partner organizations to bring products to market.

Contoso products are manufactured by using blueprint files that the company authors and maintains.

Existing Environment

Currently, Contoso uses multiple types of servers for business operations, including the following:

File servers

Domain controllers

Microsoft SQL Server servers

Your network contains an Active Directory forest named contoso.com. All servers and client computers are joined to Active Directory.

You have a public-facing application named App1. App1 is comprised of the following three tiers:

A SQL database

A web front end

A processing middle tier

Each tier is comprised of five virtual machines. Users access the web front end by using HTTPS only.

Requirements

Planned Changes

Contoso plans to implement the following changes to the infrastructure:

Move all the tiers of App1 to Azure.

Move the existing product blueprint files to Azure Blob storage.

Create a hybrid directory to support an upcoming Microsoft Office 365 migration project.

Technical Requirements

Contoso must meet the following technical requirements:

Move all the virtual machines for App1 to Azure.

Minimize the number of open ports between the App1 tiers.

Ensure that all the virtual machines for App1 are protected by backups.

Copy the blueprint files to Azure over the Internet.

Ensure that the blueprint files are stored in the archive storage tier.

Ensure that partner access to the blueprint files is secured and temporary.

Prevent user passwords or hashes of passwords from being stored in Azure.

Use unmanaged standard storage for the hard disks of the virtual machines.

Ensure that when users join devices to Azure Active Directory (Azure AD), the users use a mobile phone to verify their identity.

Minimize administrative effort whenever possible.

User Requirements

Contoso identifies the following requirements for users:

Ensure that only users who are part of a group named Pilot can join devices to Azure AD.

Designate a new user named Admin1 as the service administrator of the Azure subscription.

Admin1 must receive email alerts regarding service outages.

Ensure that a new user named User3 can create network objects for the Azure subscription.

NEW QUESTION: 150

storage1□□□ Azure Storage □□□ □□□□.
□□□ 1□ □□ Scope1□□□ □□□ □□□ □□□ □□□□□.
Scope1□ □□□□ □□ □□□□ □□□ □□□□ □ □□□?
A. □□ □□□ □□
B. □□□□□
C. □□ □□ □ □□□□□ □□
D. □□□□□ □□□□
E. □□ □□, □□□□ □ □□□□
F. □□ □□, □□□□, □□□ □ □
Answer: B (LEAVE A REPLY)
"Encryption scopes enable you to manage encryption at the level of an individual blob or container."
https://learn.microsoft.com/en-us/azure/storage/blobs/encryption-scope-manage?tabs=portal

NEW QUESTION: 151
Azure AD(Azure Active Directory) Premium□ □□□□□.
Azure AD □□□□ □□□ □□ □□□□ □□□□ admin1@contoso.com□□□ □□□□ □□□□ □□□.
Azure AD□□ □□□ □□□□ □□□?
A. □□ □□□□□ □□ □□□□□.
B. □□ □□□□□ □□ □□□□□.
C. □□□ □□□□□ □□□ □□□□□.
D. MFA □□ □□□□□ □□□□□□.
Answer: A (LEAVE A REPLY)
When you configure devices to join Azure Active Directory (Azure AD) - now called Microsoft Entra ID - you can control which users or groups are automatically assigned local administrator privileges on those devices. This configuration is done through the Device settings under the Devices blade in the Azure portal.
According to the Microsoft Azure Administrator official documentation, when a Windows 10 or later device is joined to Azure AD, the following accounts are automatically added to the local administrators group on the device:
The Azure AD global administrators.
The Azure AD device administrators (also known as "Additional local administrators on Azure AD joined devices").
The user performing the Azure AD join operation.
To assign a specific user (e.g., admin1@contoso.com) as a local administrator on all Azure AD-joined computers, an administrator must configure the Device administrators group via:
Azure Active Directory # Devices # Device settings # Additional local administrators on Azure AD joined devices.
This setting grants the selected users local administrator rights on all current and future Azure AD-joined devices within the tenant. It eliminates the need to manually configure each machine, ensuring consistent and centralized control of administrative privileges across all enrolled endpoints.
This feature is part of Azure AD Premium licensing and aligns with the Azure AD Join management policies detailed in the Azure Administrator Associate (AZ-104) study materials.

NEW QUESTION: 152
App1□□□ Azure App Services □□□ □□□□.

Web Deploy□ □□□□ App1□ □□□ □□□□□.
App1 □□□□ Azure Active Directory(Azure AD) □□ □□□ □□□□ App1□ □□□□ □□□ □ □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.
□□□ □□ □□□?
A. FTPS□ □□ □ □□ □□ □□□ □□□□□.
B. □□□□□ □□□□ □□□ □□□ □□□□□.
C. □□□□□ □□□ □□□ □□□□□.
D. FTPS□ □□ □□□ □□ □□ □□□ □□□□□.
Answer: B (LEAVE A REPLY)
When deploying to Azure App Service, several authentication methods exist:
* App-level FTPS credentials: Common to all users, not tied to Azure AD.
* User-level FTPS credentials: Separate per user, but still not Azure AD integrated.
* Azure Role-Based Access Control (RBAC): Allows assigning permissions through Azure AD roles.
To allow developers to deploy code to an App Service (via Web Deploy, Git, or Visual Studio), the least privilege role that grants this capability is the Website Contributor role.
Role capabilities:
* The Website Contributor role allows a user to manage web apps (including deployment, restart, and configuration changes), but not delete or grant access.
* It uses Azure AD authentication through RBAC, satisfying the requirement that developers use their Azure AD credentials.
Other options:
* Owner: Grants excessive permissions (manage access and delete resources) #
* App-level/User-level FTPS credentials: Use username/password, not Azure AD identities # Therefore, the correct and least-privileged role for developers to deploy via Web Deploy is Website Contributor.
# Final Verified Answer: B. Assign the Website Contributor role to the developers

NEW QUESTION: 153
VNet1□□□ Azure □□ □□□□□ □□, □ □□□□□□ Subnet1□□□ □□□□ □□□□. Subnet1□□ Azure □□ □□ □ □□ □□□, □ □□ □□□□ □□ IP □□□ □□□□.
□□ □□□ □□□ □□□□ □□ 443□ □□ □□□ □ □□ □□ □□□□□□□ □□□□□□.
□□□□□ □□□□□□ VNet1□ □□ □□□ □ VPN □□□ □□□□.
□□□□ □□□□□ □□□□□□ □□ □□□□ □□□□(RDP)□ □□□□ □□ □□□ □□□□ □ □□□ □□ □□ □□□□□.
□□□□□ □□□□□□ RDP □□□ □□□ □□□ □□□□, □□□□□ □□ □□□ □□ RDP □□□□ □□□□ □□□. □□□□ □□□ □□□□ □□ □□□□□□□ □□ □□□□ □ □□□ □□□□ □□□.
□□□ □□ □□□?
A. □□ □□□□ □□□□□□ □□ □□□ □□□□□.
B. □□ □□□□ □□ IP □□□ □□□□□.
C. Subnet1□ □□ □□□ □□□□□.
D. Subnet1□ □□□ □□□□ □□ □□(NSG)□□ □□ □□□ □□□□.
Answer: D (LEAVE A REPLY)
You can filter network traffic to and from Azure resources in an Azure virtual network with a network security group. A network security group contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure resources.
You can use a site-to-site VPN to connect your on-premises network to an Azure virtual network. Users on your on-premises network connect by using the RDP or SSH protocol over the site-to-site VPN connection.
You don't have to allow direct RDP or SSH access over the internet. And this can be achieved by configuring a deny rule in a network security group (NSG) that is linked to Subnet1 for RDP / SSH protocol coming from internet.
Modify the address space of Subnet1 : Incorrect choice
Modifying the address space of Subnet1 will have no impact on RDP traffic flow to the virtual network.
Modify the address space of the local network gateway : Incorrect choice Modifying the address space of the local network gateway will have no impact on RDP traffic flow to the virtual network.

Remove the public IP addresses from the virtual machines : Incorrect choice If you remove the public IP addresses from the virtual machines, none of the applications be accessible publicly by the Internet users.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-network/security-overview
https://docs.microsoft.com/en-us/azure/security/fundamentals/network-best-practices

NEW QUESTION: 154
ASP1□□□ Azure App Service □□□ □□□□.
ASP1□ CPU □□□□ □□ □□ □□ □□□□.

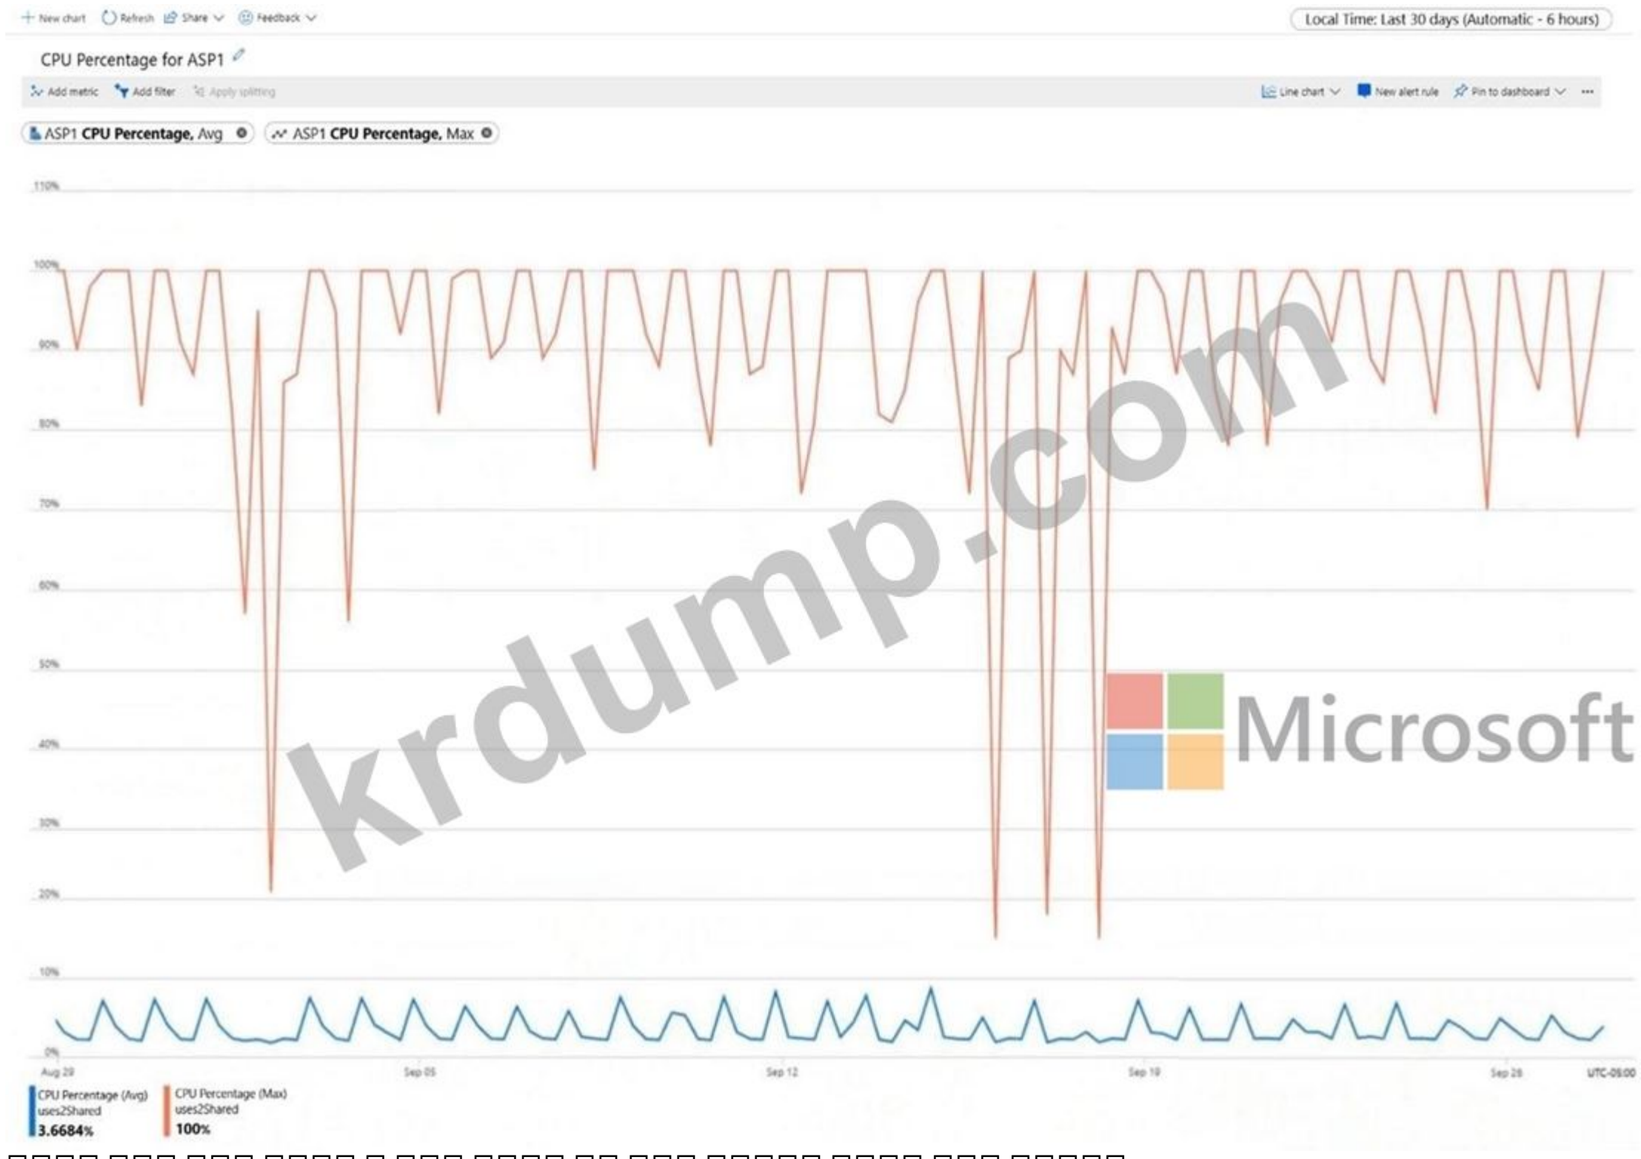

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer:

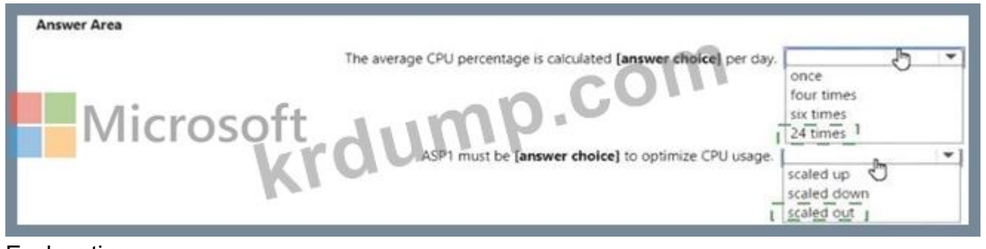

Explanation:

The average CPU percentage is calculated 24 times per day. This is because the exhibit shows the CPU percentage for ASP1 in a 24-hour period, with one data point for each hour. Therefore, the average CPU percentage is calculated once per hour, or 24 times per day1.

ASP1 must be scaled out to optimize CPU usage. This is because the exhibit shows that the CPU percentage for ASP1 is consistently above 80%, which indicates that the app service plan is under high load and needs more instances to handle the traffic. Scaling out means adding more instances to an app service plan, which can improve the performance and availability of the apps hosted on it2. Scaling up means changing the pricing tier of an app service plan, which can increase the resources available for each instance, but not necessarily reduce the CPU usage3.

NEW QUESTION: 155

□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Azure region	Resource group
VNET1	West US	RG1
VNET2	Central US	RG1
VNET3	Central US	RG2
VNET4	West US	RG2

□□ □□ Azure □□□ RG1□ AF1□□□ Azure □□□□ □□□□ □□□.

□□ □□ □□□□□ AF1□ □□□ □ □□□?

A. VNET1□

B. VNET1 □ VNET2□ □□

C. VNET1 □ VNET4□ □□

D. VNET1, VNET2 □ VNET4□ □□

E. VNET1, VNET2, VNET3 □ VNET4

Answer: C (LEAVE A REPLY)

Azure Firewall must be deployed in the same Azure region as the virtual network and into a dedicated subnet named AzureFirewallSubnet.

From the table:

* VNET1 # West US # RG1 #
* VNET4 # West US # RG2 #

* VNET2 # Central US #
* VNET3 # Central US #
Resource group location does not restrict deployment-only region alignment matters.
Microsoft documentation confirms:
"Azure Firewall must be deployed into a virtual network in the same region."

NEW QUESTION: 156
VNet1□□□ □□ □□□□□ □□□□ Subscription1□□□ Azure □□□ □□□□. □□ □□ □□□□ □□□□□.

User	Role
User1	Owner
User2	Security Admin
User3	Network Contributor

□ □□□ □□□ □ □□ □□□□ □□□□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

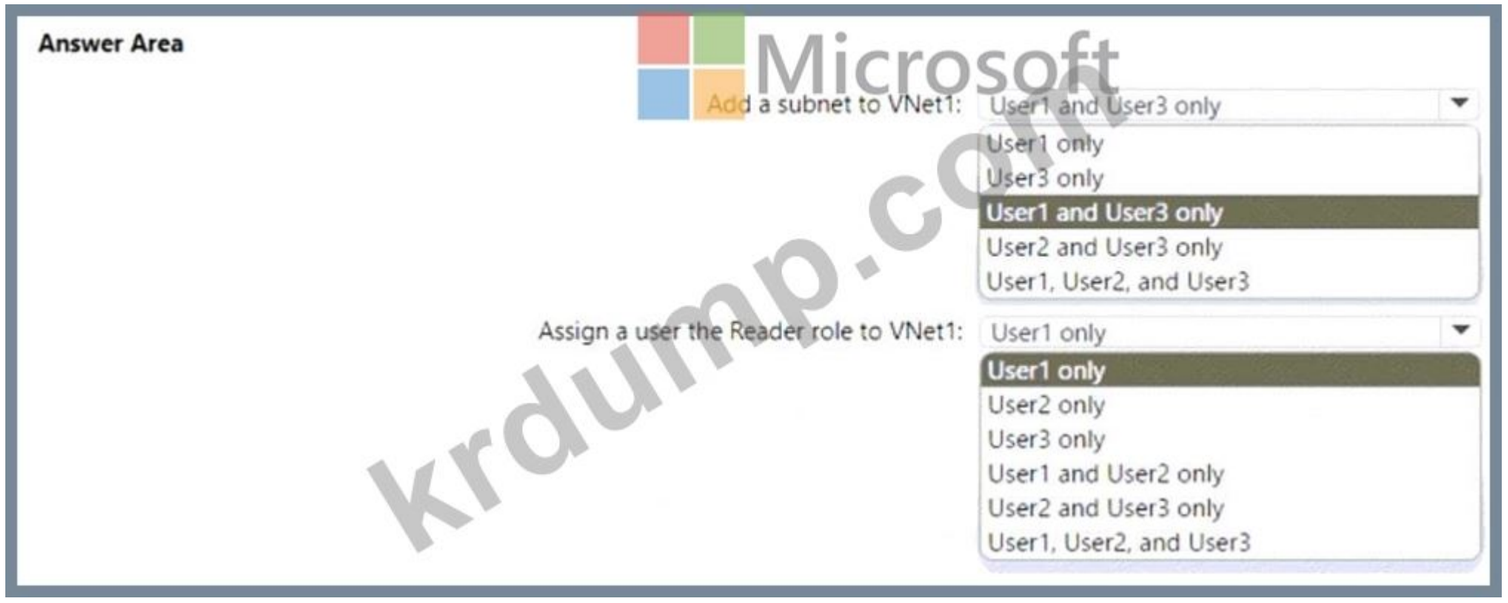

Answer:

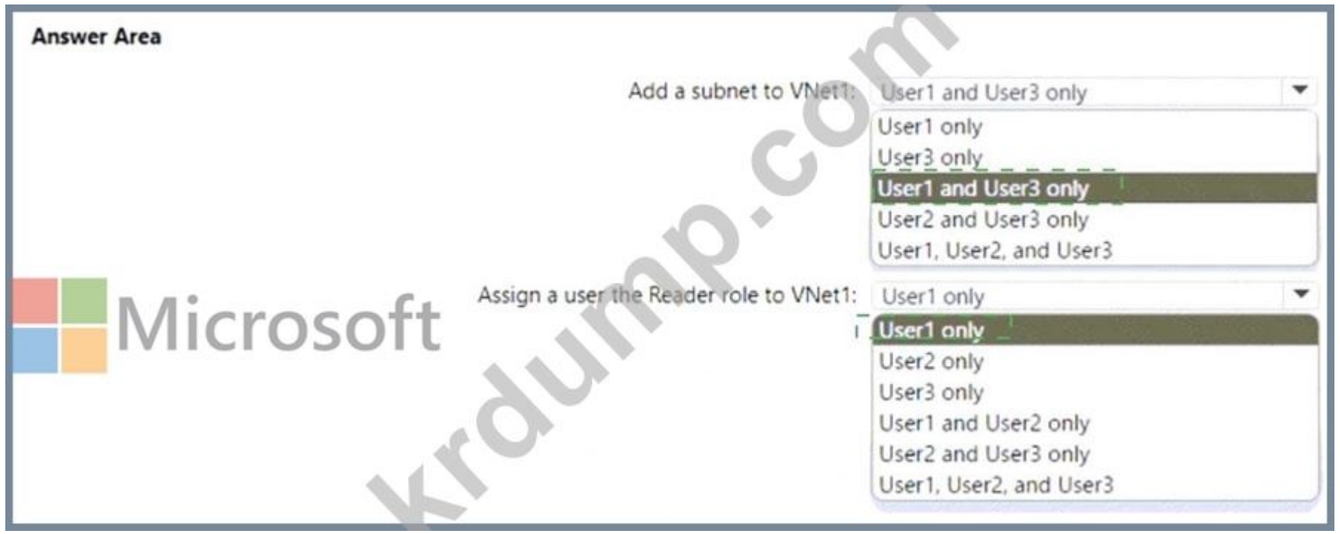

Explanation:

This question evaluates understanding of Azure Role-Based Access Control (RBAC) permissions for managing virtual networks (VNets) and assigning Azure roles.

1. Understanding Each User's Role

User

Role

Key Capabilities

User1

Owner

Full access to all resources, including the ability to delegate access by assigning roles.

User2

Security Admin

Can manage security policies and settings (Azure Defender, Security Center) but cannot manage or configure networking resources.

User3

Network Contributor

Can manage networking resources, including VNets, subnets, and network interfaces, but cannot grant access or assign roles.

2. Task Analysis

Task 1: Add a subnet to VNet1

To add a subnet to a virtual network, the user must have permission to modify network resources.

The relevant Azure RBAC action is:

Microsoft.Network/virtualNetworks/subnets/write
Roles that include this permission:
Owner - has all permissions.
Network Contributor - can manage all network settings, including adding/removing subnets.
Therefore, User1 (Owner) and User3 (Network Contributor) can perform this task.
User2 (Security Admin) cannot, because that role does not include network configuration permissions.
Task 2: Assign a user the Reader role to VNet1
To assign roles or manage access permissions, the user must have RBAC management privileges, specifically:
Microsoft.Authorization/roleAssignments/write
Only users with roles that include role assignment permissions can do this - such as:
Owner
User Access Administrator
The Network Contributor and Security Admin roles do not have this privilege.
Therefore, User1 (Owner) can assign the Reader role.
User2 and User3 cannot assign RBAC permissions.
User1 only
Microsoft Official Documentation Extracts (AZ-104 Study Guide and Azure Docs):
"The Owner role has full access to all resources, including the right to delegate access to others."
"The Network Contributor role allows management of networking resources but not access control assignments."
"The Security Admin role manages security policies and alerts but cannot create, modify, or delete network resources." (Source: Microsoft Learn # Azure built-in roles for RBAC)
Final Verified Answers Summary
Configuration
Explanation:
Add a subnet to VNet1
User1 and User3 only
Owner and Network Contributor can modify virtual network configurations.
Assign a user the Reader role to VNet1
User1 only
Only Owner can assign roles (has roleAssignments/write permission).
Therefore, the final verified answer is:
Add a subnet to VNet1 # User1 and User3 only
Assign Reader role to VNet1 # User1 only

NEW QUESTION: 157
Azure □□□ □□□□.
□□□□ □□□ □□ □□□□□ □□□ □□□□ □□□□□□. □□□ □□□□ □□ □ □□□ □ VPN□ □□□□ Azure □□□□ □□□□□ □□□. □□ □□□□ □□□□ □□□ □ VPN□ □□□□ Azure □□□□ □□□□□□.
□□ Azure □□ □□□□ □□□□ App1□□□ LOB(□□ □□) □□ □□□□. □□ □□□ Windows Server 2016□ □□□□□.
App1□ □□ □□□ □□ □□ □□□ □□□□ □□□ □□□□ □□□.
□□□ □ □□ □ □□ Azure □□□□ □□□□□? □ □□□ □□□ □□□□ □□□□□.
□□: □□□ □□□ □□ 1□□ □□□ □□□□.

A. □□ □□ □□□
B. □□□ □□□
C. Azure CDN(Content Delivery Network)
D. □□ □□ □□□
E. Azure □□□□□□ □□□□□
Answer: D,E (LEAVE A REPLY)
Line of Business WebAPP works on VMs need internal load balancer. So D is needed. Then deploy WebAPP on VMs, check the link. https://docs.microsoft.com/en-us/azure/application-gateway/quick-create-portal So B is needed as well. The orignal answer is not accomplished.

NEW QUESTION: 158
□□□ □□ VNet1□□□ □□ □□□□□ □□□□. ('□□' □□ □□□□□.)

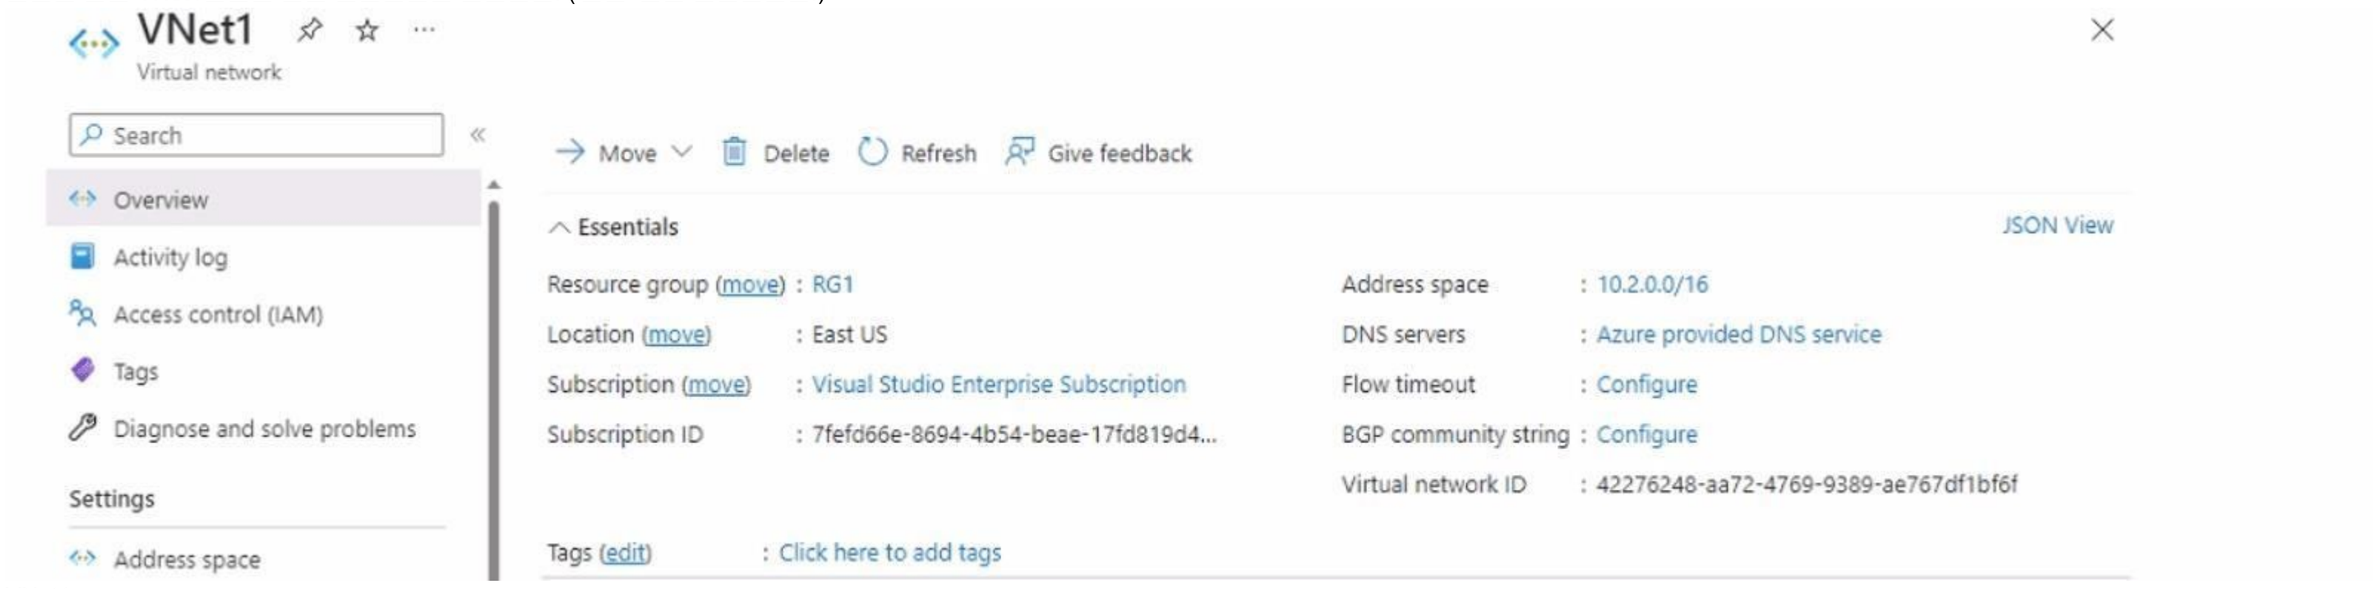

VNet1□ □□□ □□□ □□□□.
VNet1□ VNet2□□ □□ □□ □□□□□ □□□□□□ □□□. VNet2□ □□ □□□ 10.2.0.0/16□□□.
□□□□ □□□□ □□□.
□□ □□ □□□ □□ □□□?
A. VNet2□□ □□□ □□□□□□ □□□□□.
B. VNet1□ □□ □□□ □□□□□.
C. VNet1□ □□□□□ □□□□ □□□□□.
D. VNet1□ VNet2□ □□□□ □□□□.
Answer: B (LEAVE A REPLY)
To create a peering between two virtual networks, the address spaces of the virtual networks must not overlap.
VNet1 has an address space of 10.0.0.0/16, which overlaps with VNet2's address space of 10.2.0.0/16.
Therefore, you need to modify the address space of VNet1 to a non-overlapping range, such as 10.1.0.0/16, before you can create the peering. You do not need to configure a service endpoint, add a gateway subnet, or create a subnet on either virtual network for the peering to work. Then, References: [Virtual network peering]
[Modify a virtual network's address space]

NEW QUESTION: 159

□□ □□□ □□ Azure Storage □□□ □□□□.

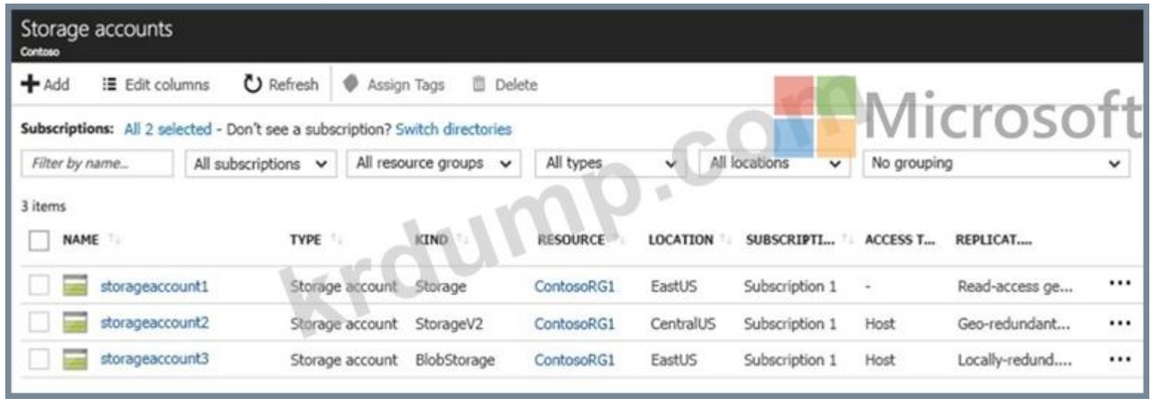

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

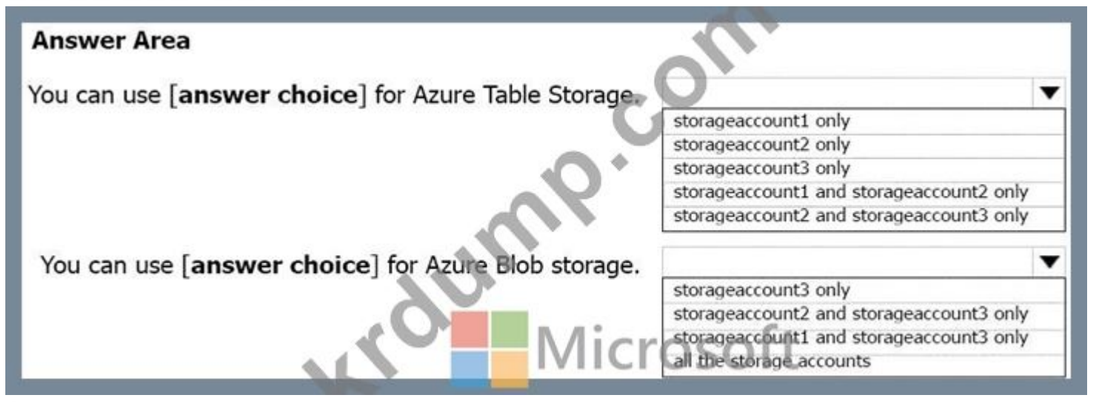

Answer:

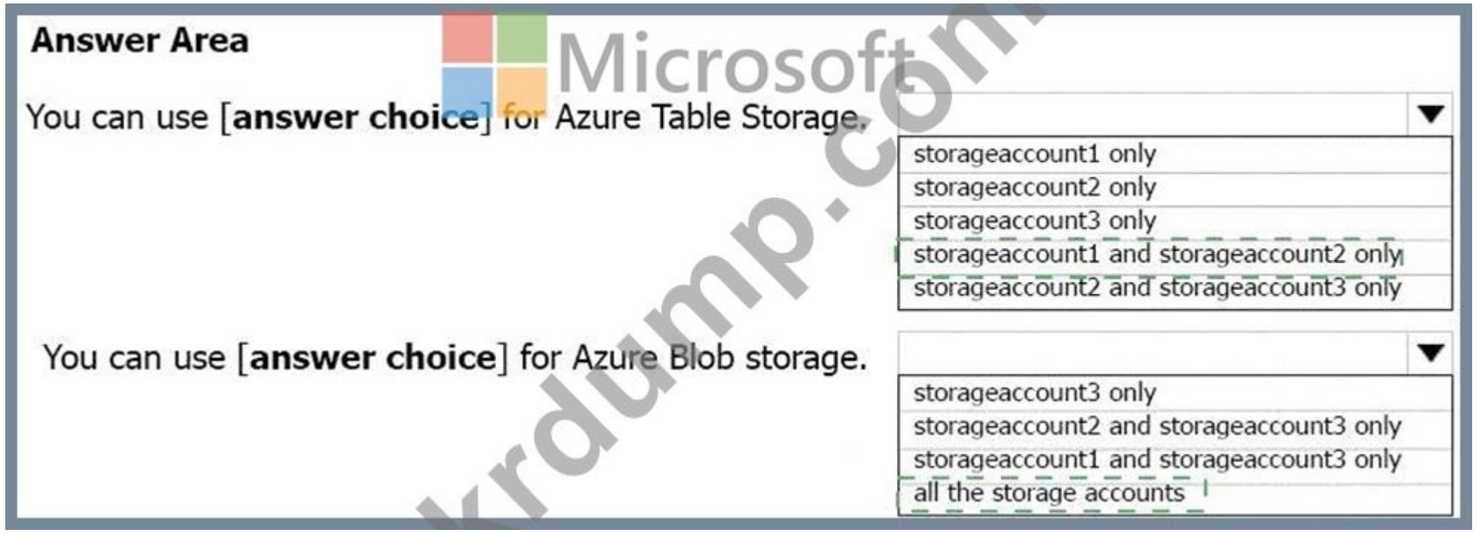

Explanation:

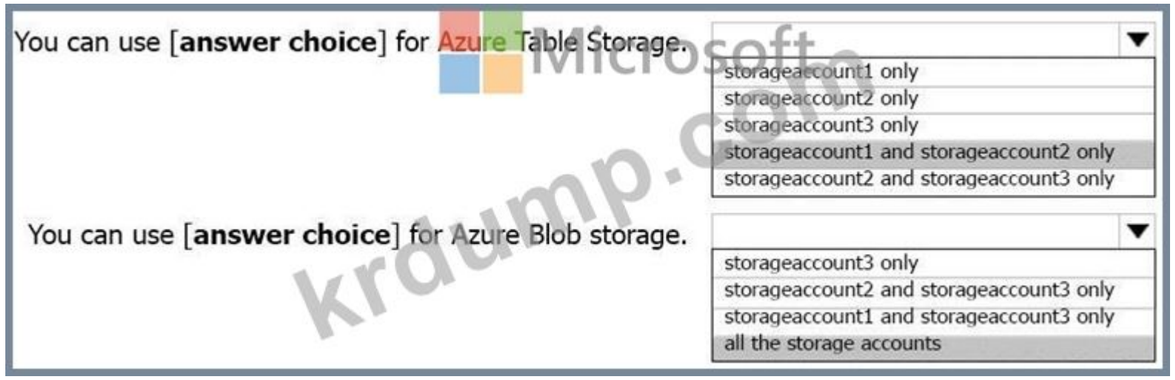

Box 1: storageaccount1 and storageaccount2 only

Box 2: All the storage accounts

Note: The three different storage account options are: General-purpose v2 (GPv2) accounts, General-purpose v1 (GPv1) accounts, and Blob storage accounts.

General-purpose v2 (GPv2) accounts are storage accounts that support all of the latest features for blobs, files, queues, and tables.

Blob storage accounts support all the same block blob features as GPv2, but are limited to supporting only block blobs.

General-purpose v1 (GPv1) accounts provide access to all Azure Storage services, but may not have the latest features or the lowest per gigabyte pricing.

References: https://docs.microsoft.com/en-us/azure/storage/common/storage-account-options

NEW QUESTION: 160

□□□ □□ □□□ □□□ □□ □□□ □□□□□ □□ □□□ □□□□ □□□.

□□ □ □□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

Answer Area

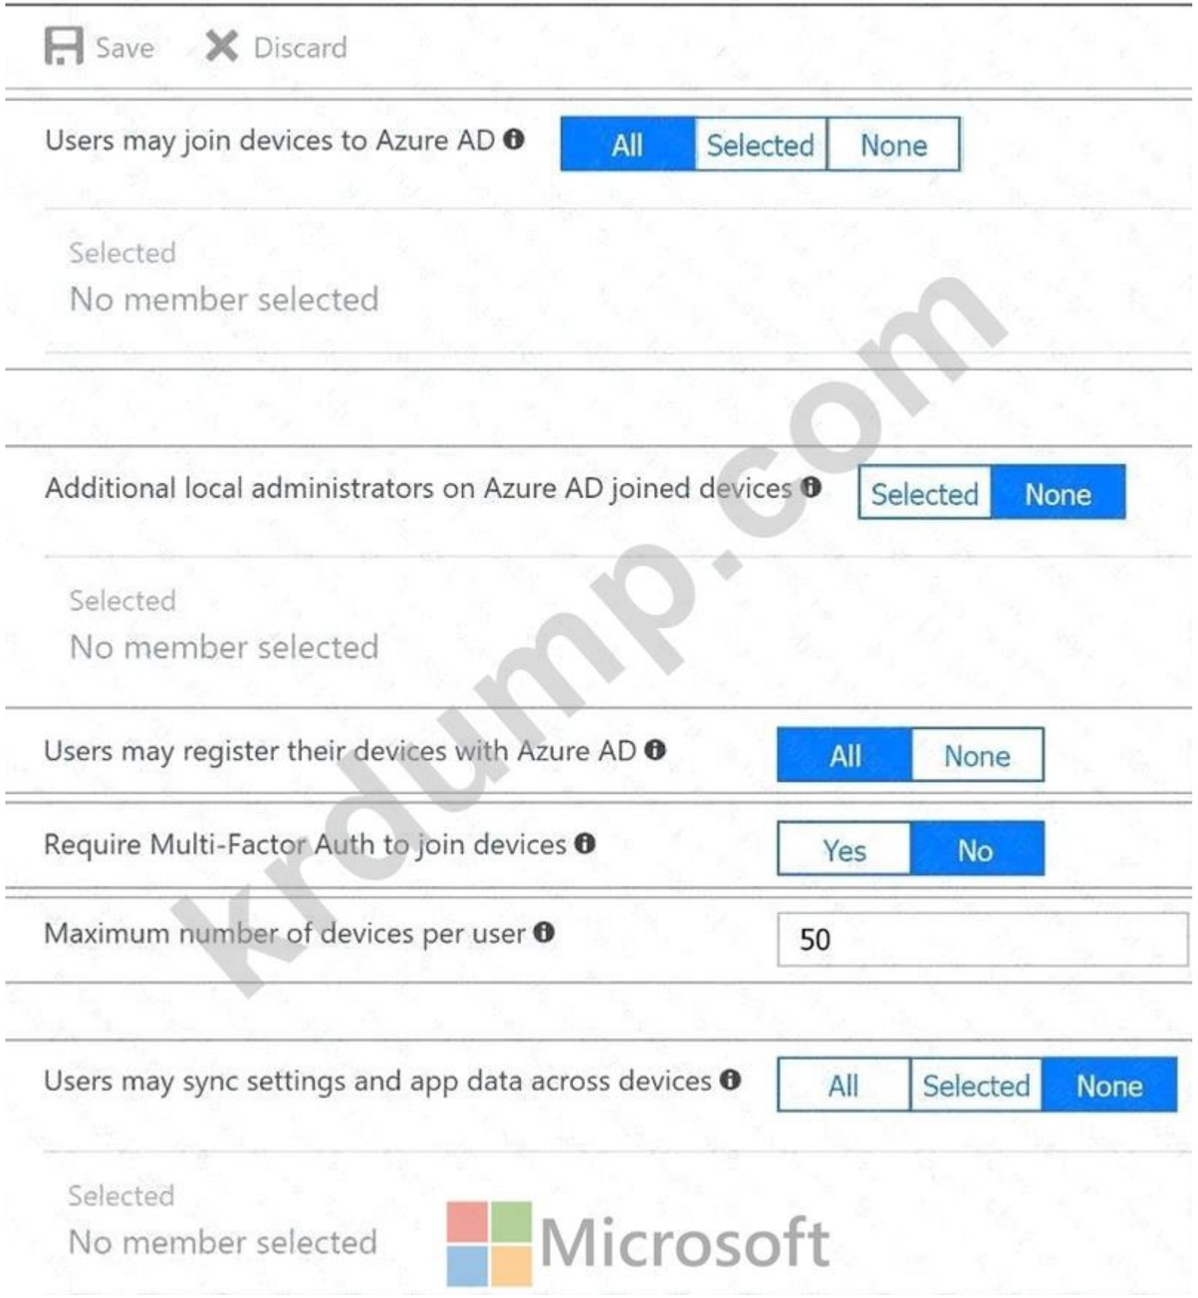

Answer:

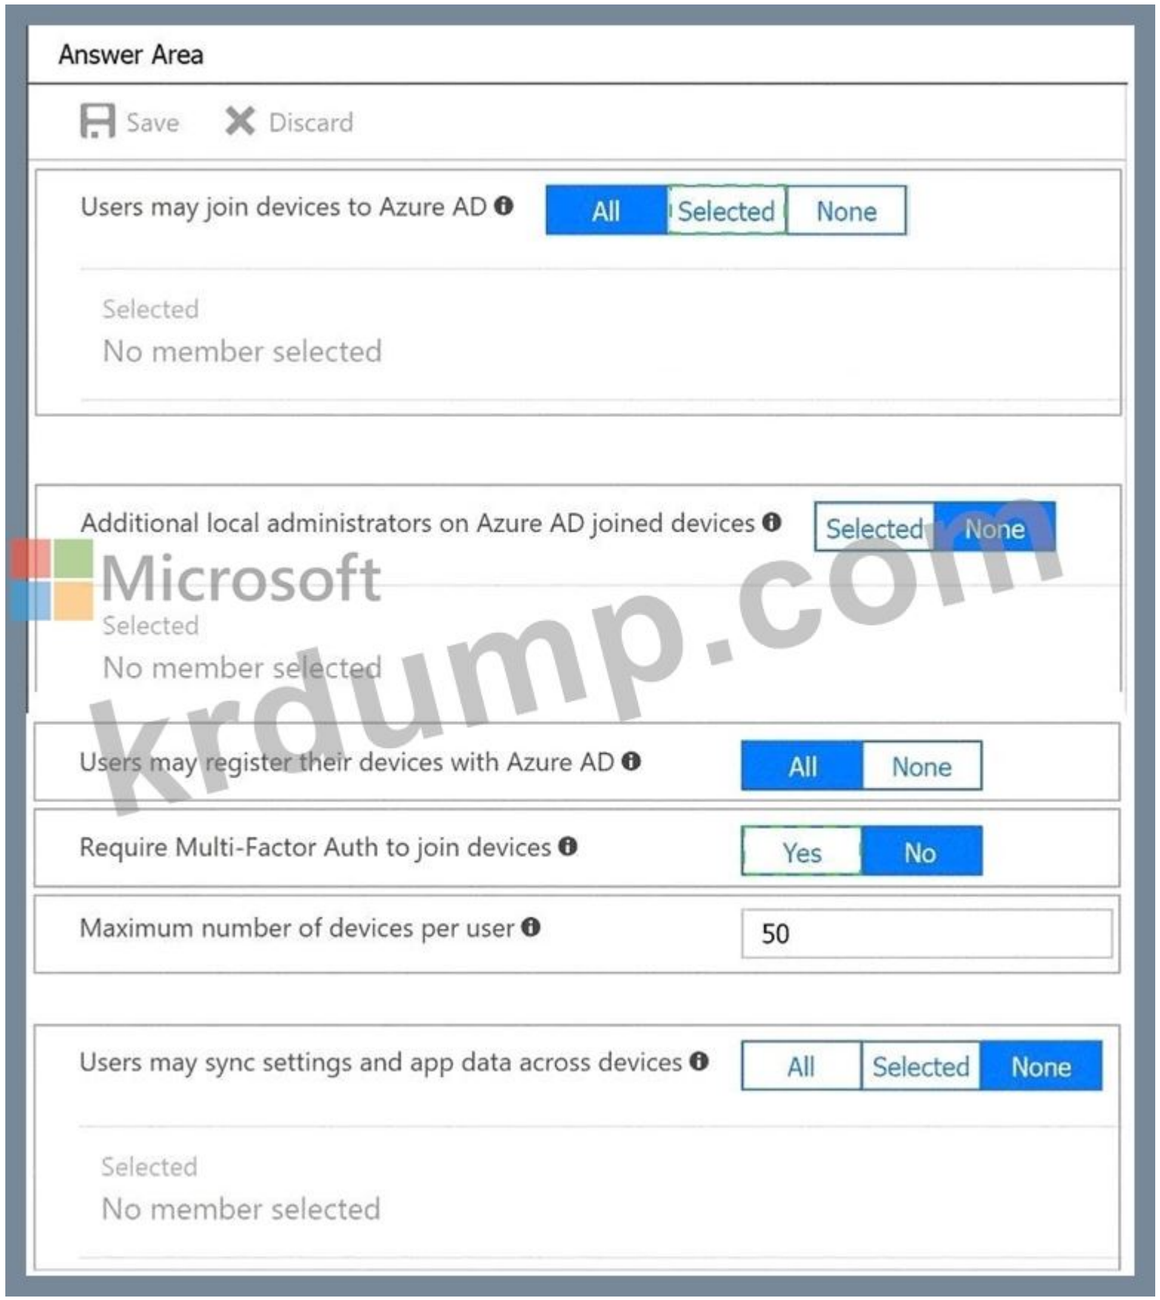

Explanation:

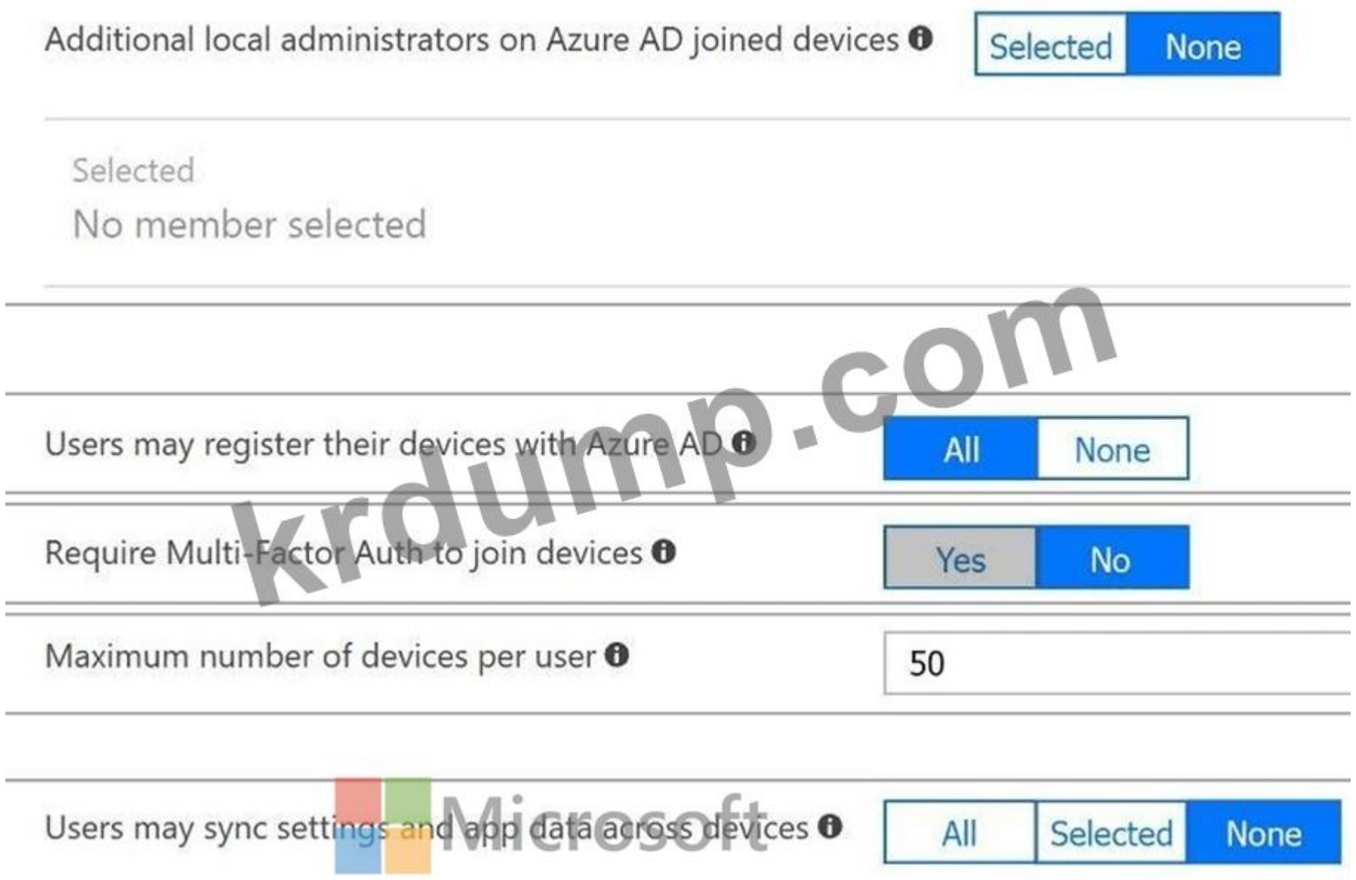

Box 1: Selected
Only selected users should be able to join devices
Box 2: Yes
Require Multi-Factor Auth to join devices.
From scenario:
Ensure that only users who are part of a group named Pilot can join devices to Azure AD Ensure that when users join devices to Azure Active Directory (Azure AD), the users use a mobile phone to verify their identity.

NEW QUESTION: 161

□□ □□ □□□ □□□ □ □□ □□□ □□□□ □□□.
□□ □ □□ □□□ □□□ □□□? □ □□□ □□□□ □□□ □□□□□.
□□: □□ 1□□ 1□□□□.
A. Azure Active Directory(AD) ID □□ □ Azure □□
B. □□ □□□ □□ □ □□ □□
C. Azure Key Vault □ □□□ □□
D. Azure Storage □□ □ □□□ □□
Answer: C (LEAVE A REPLY)
D: Seamless SSO works with any method of cloud authentication - Password Hash Synchronization or Pass- through Authentication, and can be enabled via Azure AD Connect.
B: You can gradually roll out Seamless SSO to your users. You start by adding the following Azure AD URL to all or selected users' Intranet zone settings by using Group Policy in Active Directory: https://autologon.microsoftazuread-sso.com

NEW QUESTION: 162
□□ □□ □□□ □□□ □□□ Microsoft Entra □□□□ □□□□.

Name	Type	Has an assigned license
Group1	Security	Yes
Group2	Security	No
Group3	Microsoft 365	Yes
Group4	Microsoft 365	No

□□□□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	Member of	Has a direct assigned license
User1	None	Yes
User2	Group1	No
User3	Group4	Yes
User4	None	No

□□ □□□□ □□□ □□□ □ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.
Answer Area

Users: User4 only
- User4 only
- User1 and User4 only
- User2 and User4 only
- User1, User2, User3, and User4

Groups: Group2 and Group4 only
- Group2 only
- Group2 and Group3 only
- Group2 and Group4 only
- Group1, Group2, Group3, and Group4

Answer:

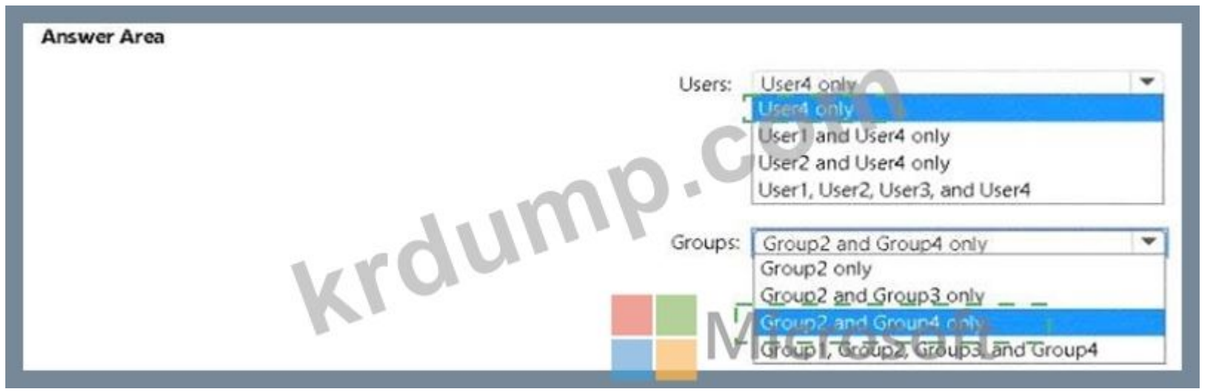

Explanation:

NEW QUESTION: 163

□□ □□ □□□ □□□ □□□ □□□ Subscription1□□□ Azure □□□ □□□□.

Name	Account kind	Azure service that contains data
storage1	Storage	File
storage2	StorageV2 (general purpose v2)	File, Table
storage3	StorageV2 (general purpose v2)	Queue
storage4	BlobStorage	Blob

Azure Import/Export □□□□ □□□□ Subscription1□□ □□□□ □□□□□ □□□.

□□ □□□ □□□□ □□□□ □□□ □ □□□?

□□□ □□□□ □□□?

A. □□□□1

B. □□□□2

C. □□□□3

D. □□□□4

Answer: D (LEAVE A REPLY)

Azure Import/Export service supports the following of storage accounts:

Standard General Purpose v2 storage accounts (recommended for most scenarios)

Blob Storage accounts

General Purpose v1 storage accounts (both Classic or Azure Resource Manager deployments), Azure Import/Export service supports the following storage types:

Import supports Azure Blob storage and Azure File storage

Export supports Azure Blob storage. Azure Files not supported.

Only storage4 can be exported.

Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-import-export-requirements

NEW QUESTION: 164
□□□ □□□□□ □□ □□□ □□□□□ □□□□ □□□?
A. □□□ □□
B. □□
C. □□□□
D. □□
Answer: (SHOW ANSWER)
You can opt in and configure additional recipients to receive your Azure invoice in an email. This feature may not be available for certain subscriptions such as support offers, Enterprise Agreements, or Azure in Open.
Select your subscription from the Subscriptions page. Opt-in for each subscription you own. Click Invoices then Email my invoice.
Click Opt in and accept the terms.
Scenario: During the testing phase, auditors in the finance department must be able to review all Azure costs from the past week.
References: https://docs.microsoft.com/en-us/azure/billing/billing-download-azure-invoice-daily-usage-date

NEW QUESTION: 165
□□ □□□ □□□□ VNetl□□□ Azure □□ □□□□□ □□□□.
* IPv4 □□ □□: 172.16.10.0/24
* □□□ □□: Subnetl
* □□□ □□ □□: 172.16.10.0/25
Subnetl□ □□□ □ □□ □□ □□□ □□ □□ □□□□□?
A. 24
B. 25
C. 123
D. 128
E. 251
Answer: (SHOW ANSWER)
In Azure Virtual Networks (VNets), each subnet defines a range of IP addresses using CIDR (Classless Inter- Domain Routing) notation. The subnet mask determines how many IP addresses are available, but Azure reserves five addresses in every subnet for its internal operations.
Given the subnet range 172.16.10.0/25, the total number of IP addresses in the subnet can be calculated as follows:
A /25 subnet provides 128 IP addresses (2^(32-25) = 128).
Azure automatically reserves five IP addresses in each subnet:
The first address (172.16.10.0) is the network identifier.
The second through fourth addresses are reserved by Azure for default gateway and internal operations.
The last address (172.16.10.127) is the broadcast address.
That leaves 128 - 5 = 123 usable IP addresses for Azure virtual machines and other resources such as NICs, load balancers, or network interfaces.
This calculation and reservation rule are explicitly documented in the Microsoft Azure Virtual Network subnet documentation. Therefore, the maximum number of virtual machines (each using a single private IP address) that can connect to the subnet is 123.

NEW QUESTION: 166

Windows Server 2019□ □□□□ VM1□□□ Azure □□ □□□ □□□□.
VM1□ Template1□□□ □□□ □□□□□ Azure Resource Manager □□□□□□ □□□□□.
Template1□□ VM2□□ □□ □□□ □□□ □□□□□.
VM2□ □□□□ □□ □□□ □□□ □ □□□?
A. □□ □□ □□
B. □□□□
C. □□□ □□□ □□
D. □□□ □□
Answer: D (LEAVE A REPLY)
Resource Group is the correct answer: Admin user, password, vm size and os are the part of ARM templates.
But resource group is not hence needs to be mentioned while deployment! Refer below sample ARM template for reference in which all above attributes passed in parameter. https://github.com/Azure/azure-quickstart-templates/blob/master/101-vm-simple-windows/azuredeploy.json

NEW QUESTION: 167
□□ Azure Resource Manager(ARM) □□□□ □□□ □□□□□.

```
{
  "$schema":"https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion":"1.0.0.0",
  "parameters":{},
  "variables":{
    "vnetId":"[resourceId('Microsoft.Network/virtualNetworks/', 'VNET1')]",
    "lbId":"[resourceId('Microsoft.Network/loadBalancers/', 'LB1')]",
    "sku":"Standard",
    "netname":"APP1"
  },
  "resources":[
    {
      "apiVersion":"2017-08-01",
      "type":"Microsoft.Network/loadBalancers",
      "name":"LB1",
      "location":"EastUS",
      "sku":{
        "name":"[variables('sku')]"
      },
      "properties":{
        "frontendIPConfigurations":[
          {
            "name":"[variables('netname')]",
  }
],
"loadBalancingRules":[
      "properties":{
      "frontendIPConfiguration":{
        "id":"[concat(variables('lbId'),'/frontendIPConfigurations/', variables('netname'))]"
      },
      "backendAddressPool":{
        "id":"[concat(variables('lbId'), '/backendAddressPools/',variables('netname'),'-Servers')]"
      },
      "probe":{
        "id":"[concat(variables('lbId'), '/probes/probe')]"
      },
      "backendPort":8080,
      "protocol":"Tcp",
      "frontendPort":80,
      "enableFloatingIP":false,
      "idleTimeoutInMinutes":4,
      "loadDistribution":"SourceIPProtocol"
    }
  }
],
"probes":[
  {
    "name":"probe",
    "properties":{
      "protocol":"Tcp",
      "port":8080,
      "intervalInSeconds":15,
      "numberOfProbes":2
    }
  }
```


□□ □ □□□ □□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

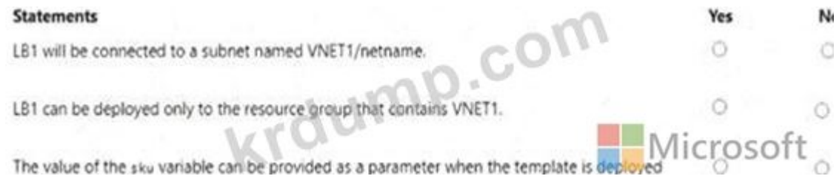

Answer:

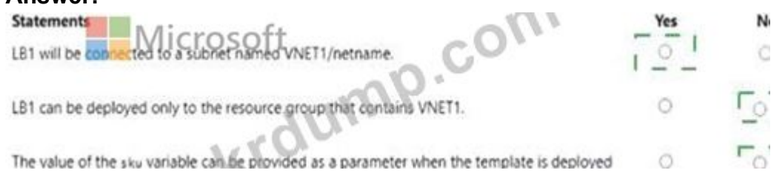

Explanation:

LB1 will be connected to a subnet named LB1 in VNET1. Yes, this is correct. The template specifies that the load balancer resource named LB1 has a property called frontendIPConfigurations, which defines the subnet where the load balancer is located. The value of this property is a reference to the resource ID of the subnet named LB1 in VNET1. You can see this reference in line 38 of the template1.

LB1 can be deployed only to the resource group that contains VNET1. No, this is not correct. The template does not specify a resource group for the load balancer resource, which means it can be deployed to any resource group in the same subscription as VNET1. However, if you want to deploy the load balancer to a specific resource group, you can add a property called resourceGroup to the reference of the subnet in line 382.

The value of the sku variable can be provided as a parameter when the template is deployed. No, this is not correct. The template defines the sku variable as a constant value of "Standard" in line 9. This means that the value cannot be changed or overridden by a parameter when the template is deployed. If you want to make the sku value configurable, you need to change the variable definition to a parameter definition, and use the parameter reference instead of the variable reference in line 363.

NEW QUESTION: 168

Admin1, Admin2, Admin3□□□ □ □□ □□□ □□□□ □□□ Azure Active Directory(Azure AD) □□□□ □□□□.

□□□□ Azure □□□ □□□□ □□□□. □□□ □□ □□□ □□□ □□□ □□ □□ □□□ □□ □□□□□. ([□] □□ □□□□□.)

Azure Portal□ Admin1□ □□□□□ □□□ □□ □□□ □□ □□□□ □□□□□. (□□ □□ □□□□□.)

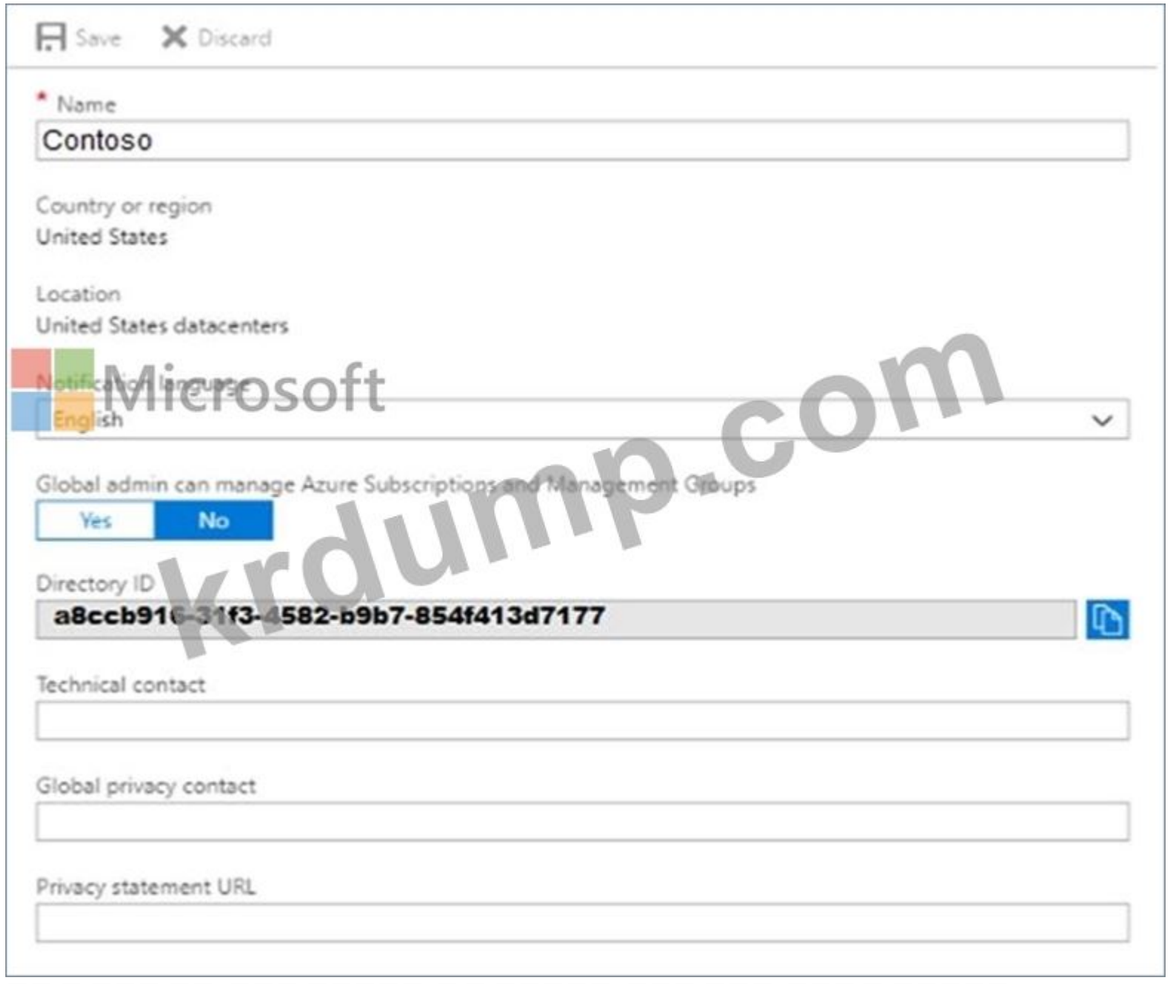

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Statements	Yes	No
Admin1 can add Admin2 as an owner of the subscription.	○	○
Admin3 can add Admin2 as an owner of the subscription.	○	○
Admin2 can create a resource group in the subscription.	○	○

Answer:

Explanation:

Explanation:

They are all Global admins so they can all modify user permission. i.e add self as owner etc.

You can be GA in one of the subscription, it doesn't mean that you can create the resources in all subscription.

As a Global Administrator in Azure Active Directory (Azure AD), you might not have access to all subscriptions and management groups in your directory. Azure AD and Azure resources are secured independently from one another. That is, Azure AD role assignments do not grant access to Azure resources, and Azure role assignments do not grant access to Azure AD.

However, if you are a Global Administrator in Azure AD, you can assign yourself access to all Azure subscriptions and management groups in your directory Reference:

https://docs.microsoft.com/en-gb/azure/role-based-access-control/elevate-access-global-admin

NEW QUESTION: 169

contoso.com□□□ Azure Active Directory(Azure AD) □□□□ □□□□.

500□□ □□ □□□□ □□□ □□□ □□□ □□□ CSV □□□ □□□□.

500□□ □□ □□□ □□□ □□ contoso.com□ □□□ □□□ □□□ □□□□ □□□.

□□ □□: □ □□□□ □□ New-MgUser cmdlet□ □□□□ PowerShell □□□□□ □□□□.

□□□ □□□ □□□□□?

A. □

B. □□□

Answer: (SHOW ANSWER)

The New-MgUser cmdlet is part of the Microsoft Graph PowerShell SDK, which is a module that allows you to interact with the Microsoft Graph API. The Microsoft Graph API is a service that provides access to data and insights across Microsoft 365, such as users, groups, mail, calendar, contacts, files, and more1.

The New-MgUser cmdlet can be used to create new users in your Azure AD tenant, but it has some limitations and requirements. For example, you need to have the Global Administrator or User Administrator role in your tenant, you need to authenticate with the Microsoft Graph API using a certificate or a client secret, and you need to specify the required parameters for the new user, such as userPrincipalName, accountEnabled, displayName, mailNickname, and passwordProfile2.

However, the New-MgUser cmdlet does not support creating guest user accounts in your Azure AD tenant.

Guest user accounts are accounts that belong to external users from other organizations or domains. Guest user accounts have limited access and permissions in your tenant, and they are typically used for collaboration or sharing purposes3.

To create guest user accounts in your Azure AD tenant, you need to use a different cmdlet: New- AzureADMSInvitation. This cmdlet is part of the Azure AD PowerShell module, which is a module that allows you to manage

your Azure AD resources and objects. The New-AzureADMSInvitation cmdlet can be used to create and send an invitation email to an external user, which contains a link to join your Azure AD tenant as a guest user. You can also specify some optional parameters for the invitation, such as the invited user display name, message info, redirect URL, or send invitation message.
Therefore, to meet the goal of creating guest user accounts for 500 external users from a CSV file, you need to use a PowerShell script that runs the New-AzureADMSInvitation cmdlet for each user, not the New-MgUser cmdlet.

NEW QUESTION: 170
App1□ □□ □□□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□. □□□□□ □□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□.

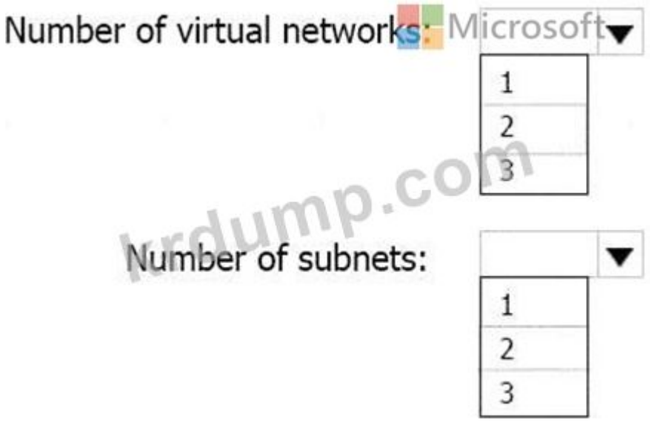

Answer:

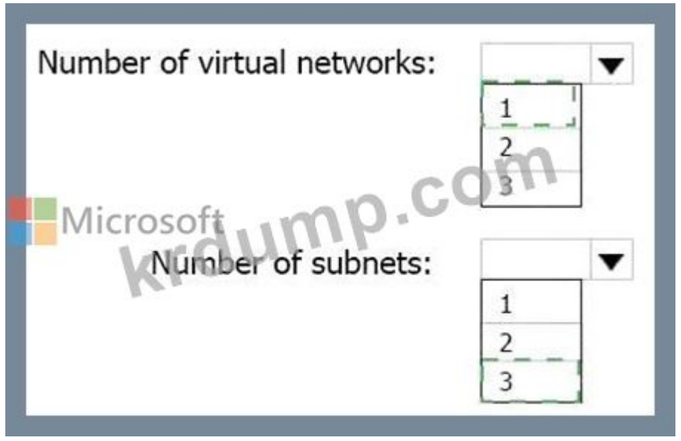

Explanation:

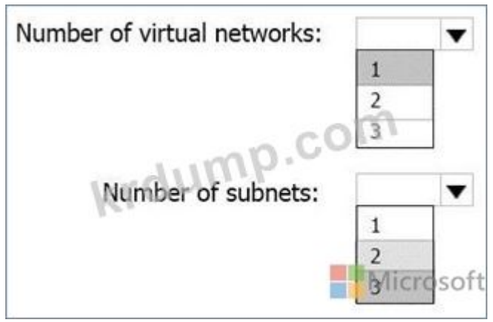

In this scenario, App1 consists of three distinct tiers - web front end, processing middle tier, and SQL database - each containing five virtual machines. The technical requirement specifies that the company must minimize the number of open ports between the App1 tiers, move all tiers of App1 to Azure, and ensure that all VMs are protected by backups.
According to Microsoft Azure architecture best practices for multi-tier applications (from Azure Architecture Center and the Azure Administrator curriculum), the optimal design involves:
* Deploying all tiers of App1 into a single virtual network (VNet).
* This allows all components of the application to communicate securely using private IP addresses.
* Keeping all tiers within a single VNet simplifies management, security, and monitoring while supporting Network Security Groups (NSGs) for inter-tier traffic control.
* Microsoft Documentation Extract:
"Use a single virtual network to host multi-tier applications. Divide the virtual network into multiple subnets, each representing a tier, and use network security groups (NSGs) to control traffic flow between tiers." (Source: Microsoft Learn - Design and implement virtual networks in Azure)
* Creating separate subnets for each application tier (3 total).
* Subnet 1: Web tier (internet-facing, HTTPS traffic)
* Subnet 2: Application/Processing tier (internal communication only)
* Subnet 3: Database tier (private, no internet access)
* Using NSGs, administrators can explicitly allow or deny traffic between subnets, thus minimizing open ports between tiers and meeting the security requirement.
* Microsoft Documentation Extract:
"Subnets provide isolation and segmentation within a virtual network. Each tier of an application should be deployed in its own subnet to apply network security policies and control exposure." (Source: Microsoft Learn - Azure virtual network design best practices)
* Backups and Storage Requirements:
* All VMs can use Azure Backup integrated with Recovery Services Vaults, which supports VM- level backup in a single VNet environment.
* The blueprint files are stored in Azure Blob Storage with the archive tier, ensuring compliance with the storage and access control requirements.
By using one virtual network and three subnets, Contoso ensures efficient management, minimized administrative overhead, secure isolation of application tiers, and full compliance with Azure governance and security recommendations.
Final Verified Answer:
* Number of virtual networks: 1
* Number of subnets: 3

NEW QUESTION: 171
Windows Server 2016□ □□□□ 5□□ Azure □□ □□□ □□□□. □□ □□□ □ □□□ □□□□□.
□□ □□□ □□ □□ □□ □□□□ □□□□ LB1□□□ Azure □□ □□ □□□ □□□□.
□ □□□ □□ □□□ □ □□□□ □□□□□ □□□□ □□□□ □□□.
□□□ □□□□ □□□?
A. □□ IP(□□ □□ □□)□ □□□□ □□

B. □□ □□ □□(□) ~ 20
C. UDP□ □□ □□□□
D. □□□□□ IP □ □□□□□ □□ □□ □□□
Answer: (SHOW ANSWER)
https://learn.microsoft.com/en-us/azure/load-balancer/distribution-mode-concepts Session persistence: Client IP and protocol - Traffic from the same client IP and protocol is routed to the same backend instance Azure Load Balancer distributes incoming network traffic among healthy backend pool instances. By default, the load balancer uses a hash-based distribution algorithm based on parameters such as source IP, destination IP, and port. This means that different requests from the same client might be directed to different backend servers.
To ensure that a client is consistently served by the same backend VM (web server) for the duration of its session, Session Persistence (also known as client affinity) must be configured.
As per the Azure Load Balancer official documentation:
"Session persistence specifies how traffic from a client should be handled by the load balancer. When you set session persistence to Client IP or Client IP and Protocol, requests from the same client IP address (and optionally protocol) are directed to the same backend instance." In web server scenarios where session state is maintained in-memory (not in a distributed cache), enabling Client IP and Protocol ensures that users remain connected to the same backend node, preventing issues like session drops or inconsistent experience.
Other options do not meet the requirement:
Floating IP (direct server return): Used for SQL Always On or NVA scenarios, not session affinity.
Idle timeout: Defines TCP connection time-out, not persistence.
UDP protocol: Not relevant for HTTP-based traffic.
Hence, the correct and Microsoft-documented answer is Session persistence = Client IP and Protocol.

NEW QUESTION: 172
Azure □□□ □□□□. □□□□ □□ □□ □□□ □□ □□ □□ □□□ □□ Storage1□□□ □□□□ □□□ □□□□ □□□□.

Name	If base blobs were last modified more than (days)	Then
Rule1	5 days	Move to cool storage
Rule2	5 days	Delete the blob
Rule3	5 days	Move to archive storage

6□ 1□□ Storage1□ □ □□□ □□□ File1□□□ Blob□ □□□□□. 6□□ File1 □□□ □□□□□?
7?
A. □□□□ □□□ □□□ □□□□□.
B. □ □□□ □□□ □□□
C. □ □□□ □□□ □□□
D. □□□
Answer: D (LEAVE A REPLY)
If you define more than one action on the same blob, lifecycle management applies the least expensive action to the blob. For example, action delete is cheaper than action tierToArchive. Action tierToArchive is cheaper than action tierToCool. https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview

NEW QUESTION: 173
□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□□ □□□□□. □□□□ □ □□□□ □□□ □□□ □□□ □ □□ □□□ □□□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □ □□ □□ □ □□□ □□ □ □ □□□□ □□□ □□ □□ □□□□.
□ □□□ □□□ □□□ □□□ □□ □□□□ □□□ □ □□□□. □□□□□ □□□ □□□ □□ □□□ □□□□ □□□□.
10□□ □□ □□□□□ □□□□ Azure □□□ □□□□. □□ □□□□□ □□□ □□□ □□□□ □□□□□□.
□□ □□□□ □□□ □□ NSG(□□□□ □□ □□)□ □□ □□□□□.

NSG□ □□□□ □□ □□□□ □□ TCP □□ 8080□ □□□□ □□□□□ □□□□ □□□.
□□ □□: □□□ □□□ □□□□□□ Microsoft.ClassicNetwork □□□□ □□ □□□□□.
□□□ □□□ □□□□□?
A. □
B. □□□
Answer: B (LEAVE A REPLY)
No, this does not meet the goal. Unregistering the Microsoft.ClassicNetwork provider does not affect the creation of network security groups (NSGs) in the subscription. The Microsoft.ClassicNetwork provider is used for managing classic deployment model resources, such as virtual networks, network interfaces, and public IP addresses1. However, NSGs are only supported for Resource Manager deployment model resources2. Therefore, unregistering the Microsoft.ClassicNetwork provider will not automatically block TCP port 8080 between the virtual networks.
To meet the goal, you need to create a custom policy definition that enforces a default security rule for NSGs. A policy definition is a set of rules and actions that Azure performs when evaluating your resources3.
You can use a policy definition to specify the required properties and values for NSGs, such as the direction, protocol, source, destination, and port of the security rule. You can then assign the policy definition to the subscription scope, so that it applies to all the resource groups and virtual networks in the subscription.

NEW QUESTION: 174
□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Subnet	Subnet-associated network security group (NSG)	Peered with
VNet1	Subnet1	NSG1	VNet2
VNet2	Subnet2	NSG2	VNet1

□□□□ □□ □□ □□□ □□ □□□ □□□□ □□□□.

Name	Connected to
VM1	Subnet1
VM2	Subnet2

□□□□ □□ □□ □□□ Azure App Service □□□ □□□□ □□□□.

Name	Description
WebApp1	Uses the Premium pricing tier and has virtual network integration with VNet1
WebApp2	Uses the Isolated pricing tier and is deployed to Subnet2

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.
Answer Area

Statements	Yes	No
WebApp1 can communicate with VM2.	○	○
NSG1 controls inbound traffic to WebApp1.	○	○
WebApp2 can communicate with VM1.	○	○

Answer:

Explanation:

WebApp1 can communicate with VM2. No, this is not correct. According to the tables, WebApp1 is integrated with VNet1, which has a peering connection with VNet2. However, VM2 is in VNet3, which is not peered with VNet1 or VNet2. Therefore, WebApp1 cannot communicate with VM2 across different virtual networks1.

NSG1 controls inbound traffic to WebApp1. No, this is not correct. According to the tables, NSG1 is associated with Subnet1 in VNet1, which is integrated with WebApp1. However, network security groups only control outbound traffic from App Service apps to virtual networks, not inbound traffic to App Service apps from virtual networks2. Therefore, NSG1 does not control inbound traffic to WebApp1.

WebApp2 can communicate with VM1. Yes, this is correct. According to the tables, WebApp2 is integrated with VNet3, which has a peering connection with VNet2. VM1 is in Subnet2 in VNet2, which has a network security group named NSG2 that allows inbound traffic from any source on port 803. Therefore, WebApp2 can communicate with VM1 on port 80 across peered virtual networks.

NEW QUESTION: 175

□□ □□ □□□ □□ □□□ □□□ □□□ Azure □□□ □□□□.

Name	Resource group	Daily cost
VM1	RG5	20 euros
VM2	RG6	30 euros

□□ □□ □□ □□ □□□ □□□ □□□.

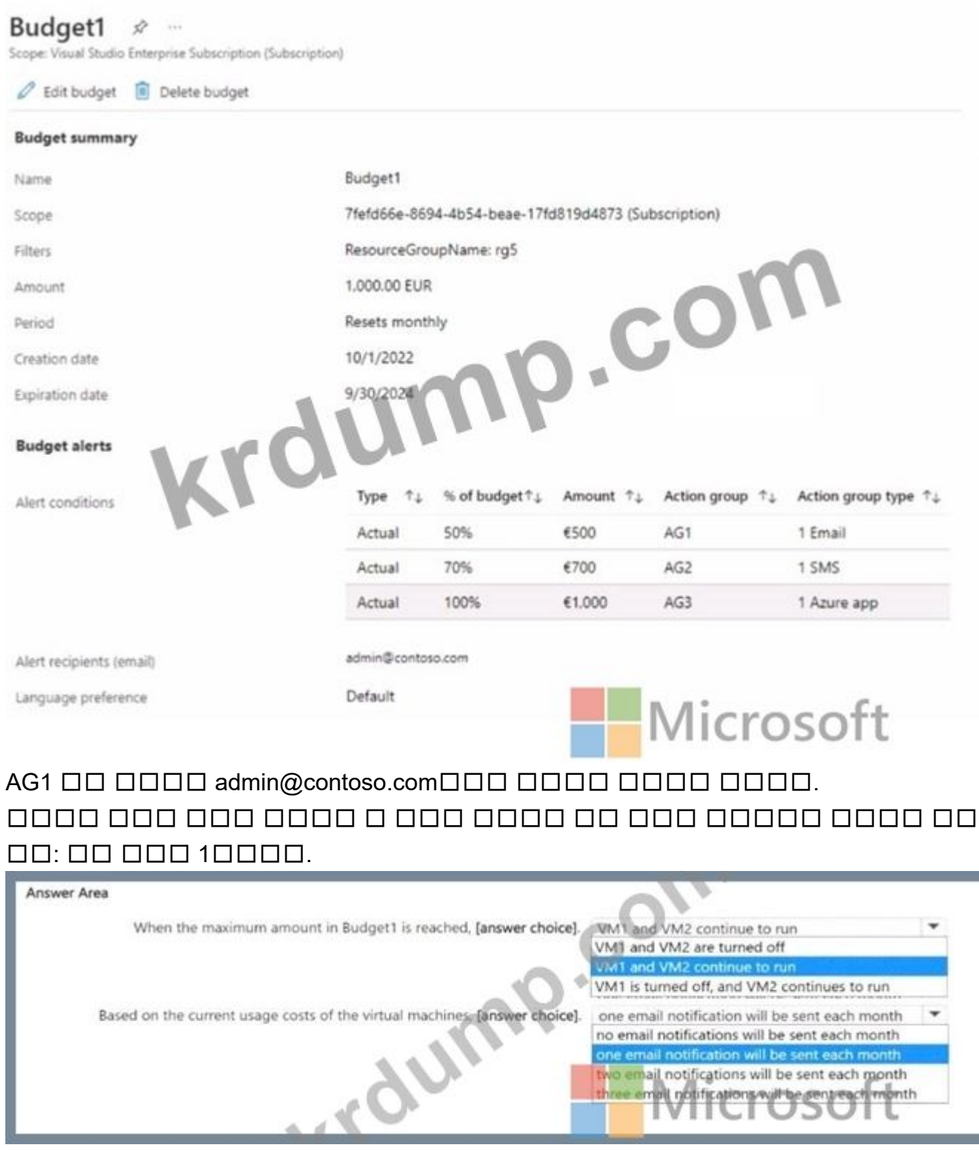

AG1 □□ □□□□ admin@contoso.com□□□ □□□□ □□□□ □□□□.

□□□□ □□□ □□□ □□□□ □ □□□ □□□□ □□ □□□ □□□□□ □□□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer:

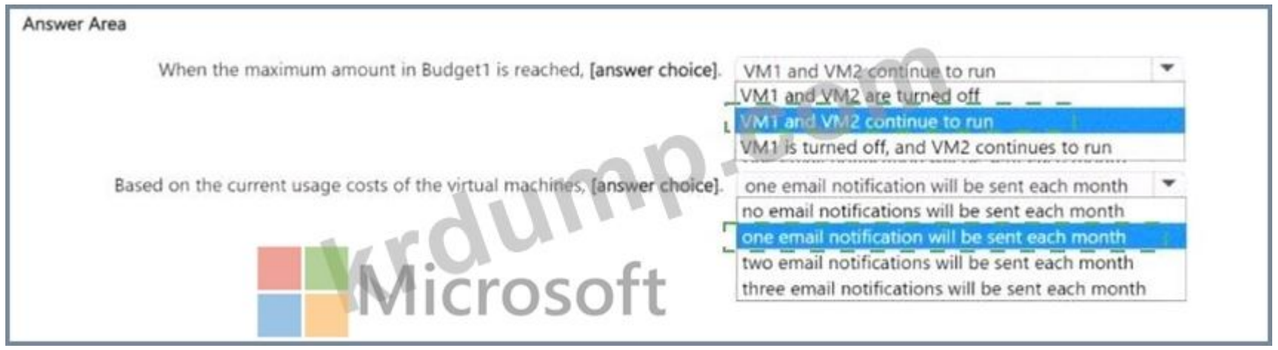

Explanation:

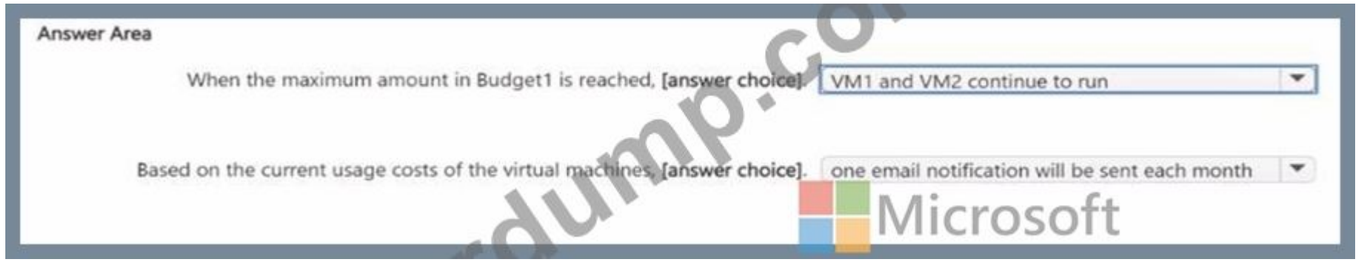

NEW QUESTION: 176

□□ □□ □□□□ □□□ Azure □□□ □□□□.

Name	Type	Details
VNet1	Virtual network	*Not applicable*
Subnet1	Subnet	Hosted on VNet1
VM1	Virtual machine	On Subnet1
VM2	Virtual machine	On Subnet1

VM1□ VM2□ □□□ □□□□□ □□□□, □□ □□□□□ □□ □□□□□ LOB(□□ □□) □□□□□□□ □□□□□□. □□□ □□□ □□□□ □□ □□(NSG)□ □□□□□. (□□ □□□ □□□□□.)

→ Move Delete

Resource group (change)
ProductionRG

Location
North Europe

Subscription (change)
Production subscription

Subscription ID
14d26092-8e42-4ea7-b770-9dcef70fb1ea

Tags (change)
Click here to add tags

Security rules
1 inbound, 1 outbound

Associated with
0 subnets, 0 network interfaces

Inbound security rules

PRIORITY	NAME	PORT	PROTOCOL	SOURCE	DESTINATION	ACTION
1500	Port_80	80	TCP	Internet	Any	Deny ...
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	AllowAzureLoadBalancerInBound	Any	Any	AzureLoadBalancer	Any	Allow ...
65500	DenyAllBound	Any	Any	Any	Any	Deny ...

Outbound security rules

PRIORITY	NAME	PORT	PROTOCOL	SOURCE	DESTINATION	ACTION
1000	DenyWebSites	80	TCP	Any	Internet	Deny ...
65000	AllowVnetOutBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow ...
65001	AllowInternetOutBound	Any	Any	Any	Internet	Allow ...
65500	DenyAllOutBound	Any	Any	Any	Any	Deny ...

VM1□ VM2 □□□□ □□□□ □□□□□ □□□□ □□ □□□□ □□□.
□□□ □□ □□□?

A. NSG□ Subnet1□ □□□□□.
B. NSG□ □□□□ □□□□□□□ □□□□□.
C. DenyWebSites □□□□□ □□ □□□ □□□□□.
D. Port_80 □□□□ □□ □□ □□

Answer: A (LEAVE A REPLY)

Outbound rule "DenyWebSites" is setup correctly to block outbound internet traffic over port 80. In the screenshot it states, "Associated with: 0 subnets, 0 NIC's", so you need to associate the NSG to Subnet1.You can associate or dissociate a network security group from a NIC or Subnet. Reference: https://docs.microsoft.

com/en-us/azure/virtual-network/manage-network-security-group

NEW QUESTION: 177

WebApp1□□□ □ Azure App Service □□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.

* WebApp1□ app.contoso.com□ □□□ □□ □□□ □□□ □□□ □ □□□ □□□.
* WebApp1□ □□ 8□□ □□□□□□ □□□□ □□□ □ □□□ □□□.
* □□□ □□□ □□□ □□□□□ □□□.

□□ □□ □□□ □□□□ □□, □□□□ □□□□ □ □□ □□□ □□□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

Answer:

Explanation:

When configuring an Azure App Service app (WebApp1), several hosting plans determine scaling capabilities and supported features such as custom domains, SSL, and auto-scaling.
According to Microsoft Azure Administrator Documentation:
* Custom Domain Verification:To map a custom domain (like app.contoso.com) to an Azure Web App, Azure must verify that you own the domain.Microsoft documentation specifies:
"To verify a custom domain, create a TXT record in your DNS zone. The TXT record contains a verification token that Azure uses to confirm ownership of the domain before binding it to your App Service." (Source: Azure App Service - Map custom domain) While an A record or CNAME record is eventually used to direct traffic to the app, the TXT record is used solely for domain verification.
* Scaling Requirement (Up to 8 Instances):The Free and Shared App Service plans have significant limitations:
* Free / Shared: No scaling (only 1 instance, no custom domain SSL support).
* Basic: Supports up to 3 instances.
* Standard: Supports up to 10 instances, includes auto-scaling, and supports custom domains and SSL.
* Premium: Adds advanced scaling and isolation but is more expensive and unnecessary here.
Microsoft documentation states:
"The Standard pricing tier supports auto-scaling up to 10 instances and custom domains, offering a balance between cost and capability." (Source: Azure App Service Plan Tiers Overview) Since the requirement includes automatic scaling up to eight instances and custom domain verification, the Standard plan is the minimum suitable tier - satisfying both performance and cost-efficiency conditions.
Final Verified Answer:
* Pricing plan: Standard
* Record type: TXT
Justification Summary (from Microsoft Documentation):
* Custom domain verification # Requires a TXT record.
* Auto-scale up to 8 instances # Requires at least the Standard plan.
* Minimize cost and effort # Standard plan is the optimal balance.

NEW QUESTION: 178
□□ □□ □□□ □□□□ □□□ contoso.onmicrosoft.com□□□ Microsoft Entra □□□□ □□□□.

Name	Member of	Role assigned
User1	Group1	None
User2	Group2	None
User3	Group1, Group2	User Administrator

□□□□ □□□ □□□ □□□ □□ contoso.onmicrosoftxom□ □□ □□□□ □□□□ □□□□□□. ('□□□□ □□□' □□ □□□□□.)

□□□□ □□□□ □□ □□ □□□ □□ □□ □□ □□□ □□ □□□□□.
(□□ □□ □□ □□□□□.)

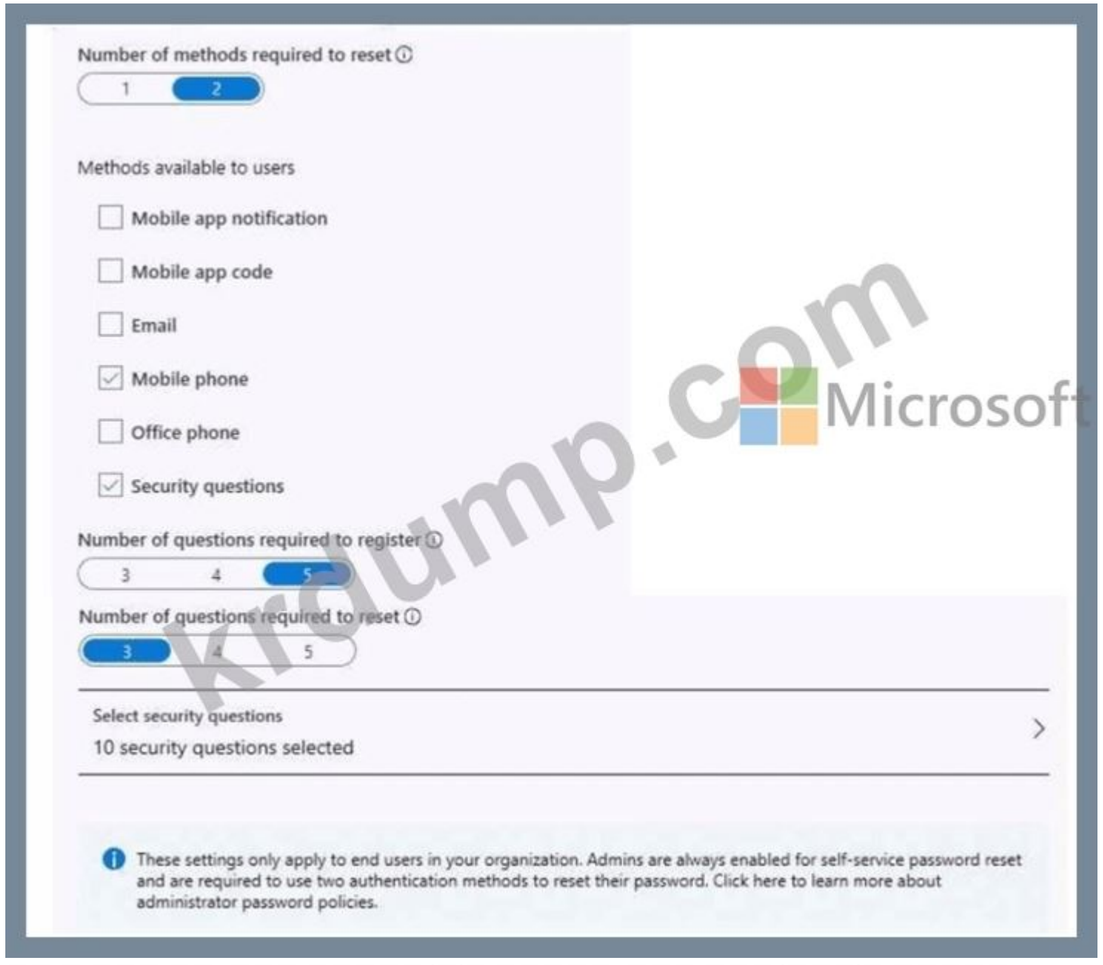

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
After User2 answers three security questions correctly, he can reset his password immediately.	○	○
If User1 forgets her password, she can reset the password by using the mobile phone app.	○	○
User3 can add security questions to the password reset process.	○	○

Answer:

Answer Area

Statements	Yes	No
After User2 answers three security questions correctly, he can reset his password immediately.	○	○
If User1 forgets her password, she can reset the password by using the mobile phone app.	○	○
User3 can add security questions to the password reset process.	○	○

Explanation:

Answer Area

Statements	Yes	No
After User2 answers three security questions correctly, he can reset his password immediately.	○	●
If User1 forgets her password, she can reset the password by using the mobile phone app.	○	●
User3 can add security questions to the password reset process.	●	○

NEW QUESTION: 179

□□ □□ □□□ □□ □□□□□ □□□ Azure □□□ □□□□.

Name	Peered with	DNS server
VNET1	VNET2	Default (Azure-provided)
VNET2	VNET1	10.10.0.4

□□ □□ □□ □□□ □□ □□□□.

Name	IP address	Network interface	Connects to
Server1	10.10.0.4	NIC1	VNET1/Subnet1
Server2	172.16.0.4	NIC2	VNET1/Subnet2
Server3	192.168.0.4	NIC3	VNET2/Subnet2

□□ □□ □□ □□□□ □□□□□□ □□ □□□□.

Name	DNS server
NIC1	Inherit from virtual network
NIC2	10.10.0.4
NIC3	Inherit from virtual network

Server1□ □□ □□ □□□ □□□□ □□□□ DNS □□□□□.

Name	Type	Value
contoso.com	Primary DNS zone	**Not applicable**
Host1.contoso.com	A record	131.107.10.15

VNET2□ □□ □□ □□□□ □□□ □□ contoso.com□□□ Azure □□ DNS □□□ □□ □□ □□ □□□□ □□ □□□□.

Name	Type	Value
Host1	A record	131.107.200.20
Host2	A record	131.107.50.50

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Statements	Yes	No
Server2 resolves host2.contoso.com to 131.107.50.50.	○	○
Server2 resolves host1.contoso.com to 131.107.10.15.	○	○
Server3 resolves host2.contoso.com to 131.107.50.50.	○	○

Answer:

Statements	Yes	No
Server2 resolves host2.contoso.com to 131.107.50.50.	○	○
Server2 resolves host1.contoso.com to 131.107.10.15.	○	○
Server3 resolves host2.contoso.com to 131.107.50.50.	○	○

Explanation:
NYN

NEW QUESTION: 180
□□ 1□ □□ 2□□ □ □□ Azure □□□ □□□□□.
□□!□□ □□□ □□ □□ □□□□□ □□□□.

Name	Address space	Region
VNET1	10.10.10.0/24	West Europe
VNET2	172.16.0.0/16	West US

□□ □□□□□□ □□□ □□ □□□□ □□□□□.

Name	Address range	In virtual network
Subnet11	10.10.10.0/24	VNET1
Subnet21	172.16.0.0/18	VNET2
Subnet22	172.16.128.0/18	VNET2

Subscnption2□□ □□□ □□ □□ □□□□□ □□□□ □□□□.
- □□: VNETA
* □□ □□: 10.10.128.0/17
* □□: □□□ □□
VNETA□□ □□□ □□ □□□□ □□□□ □□□□.

Name	Address range
SubnetA1	10.10.130.0/24
SubnetA2	10.10.131.0/24

□□ □ □□□ □□, □□□ □□□□□ '□'□ □□□□□. □□□ □□□ '□□□'□ □□□□□.
□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
A Site-to-Site connection can be established between VNET1 and VNET2.	○	○
VNET1 and VNET2 can be peered.	○	○
VNET1 and VNETA can be peered.	○	○

Answer:

Answer Area

Statements	Yes	No
A Site-to-Site connection can be established between VNET1 and VNET2.	○	○
VNET1 and VNET2 can be peered.	○	○
VNET1 and VNETA can be peered.	○	○

Explanation:

Answer Area

Statements	Yes	No
A Site-to-Site connection can be established between VNET1 and VNET2.	●	○
VNET1 and VNET2 can be peered.	●	○
VNET1 and VNETA can be peered.	○	●

In Microsoft Azure, connectivity between virtual networks depends on IP address space, region, and subscription constraints.
1## Site-to-Site VPN between VNET1 and VNET2 - YES:
Azure supports a Site-to-Site VPN connection between two VNets in different regions or subscriptions as long as both have non-overlapping address spaces and contain a VPN gateway. In this case, VNET1 (10.10.10.0/24 - West Europe) and VNET2 (172.16.0.0/16 - West US) have unique address spaces, so a S2S connection is possible.
2## VNET1 and VNET2 Peering - YES:
VNet peering enables direct communication over the Microsoft backbone network with low latency and no gateway requirement. Global VNet Peering supports VNets in different Azure regions (e.g., West Europe and West US) as long as address spaces do not overlap-which is true here.
3## VNET1 and VNETA Peering - NO:
VNETA's address space (10.10.128.0/17) overlaps with VNET1 (10.10.10.0/24), because both belong to the same 10.10.0.0/8 range. Azure explicitly blocks VNet peering between VNets with overlapping address spaces to avoid routing conflicts.
According to Microsoft Learn ("Virtual network peering - requirements and constraints"), peering is only supported when address spaces do not overlap. Therefore, VNET1 and VNETA cannot be peered.