Microsoft.AZ-104.v2026-04-16.q199

Exam Code:	AZ-104
Exam Name:	Microsoft Azure Administrator
Certification Provider:	Microsoft
Free Question Number:	199
Version:	v2026-04-16
# of views:	115
# of Questions views:	2711
https://www.freecram.net/torrent/Microsoft.AZ-104.v2026-04-16.q199.html	

NEW QUESTION: 1

You have an Azure subscription that contains the resources in the following table.

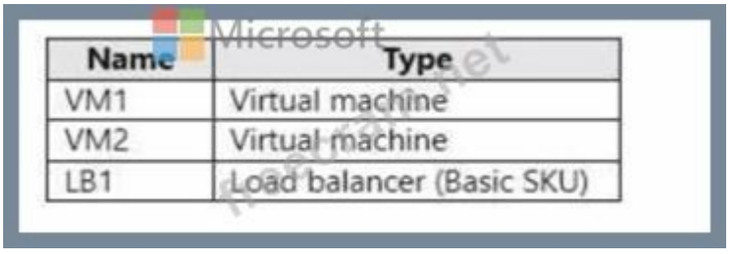

Name	Type
VM1	Virtual machine
VM2	Virtual machine
LB1	Load balancer (Basic SKU)

You install the Web Server (IIS) server role on VM1 and VM2. and then add VM1 and VM2 to LB1.

LB1 is configured as shown in the LB1 exhibit. (Click the LB1 tab.)

Rule1 is configured as shown in the Rule1 exhibit (Click the Rule1 tab.)

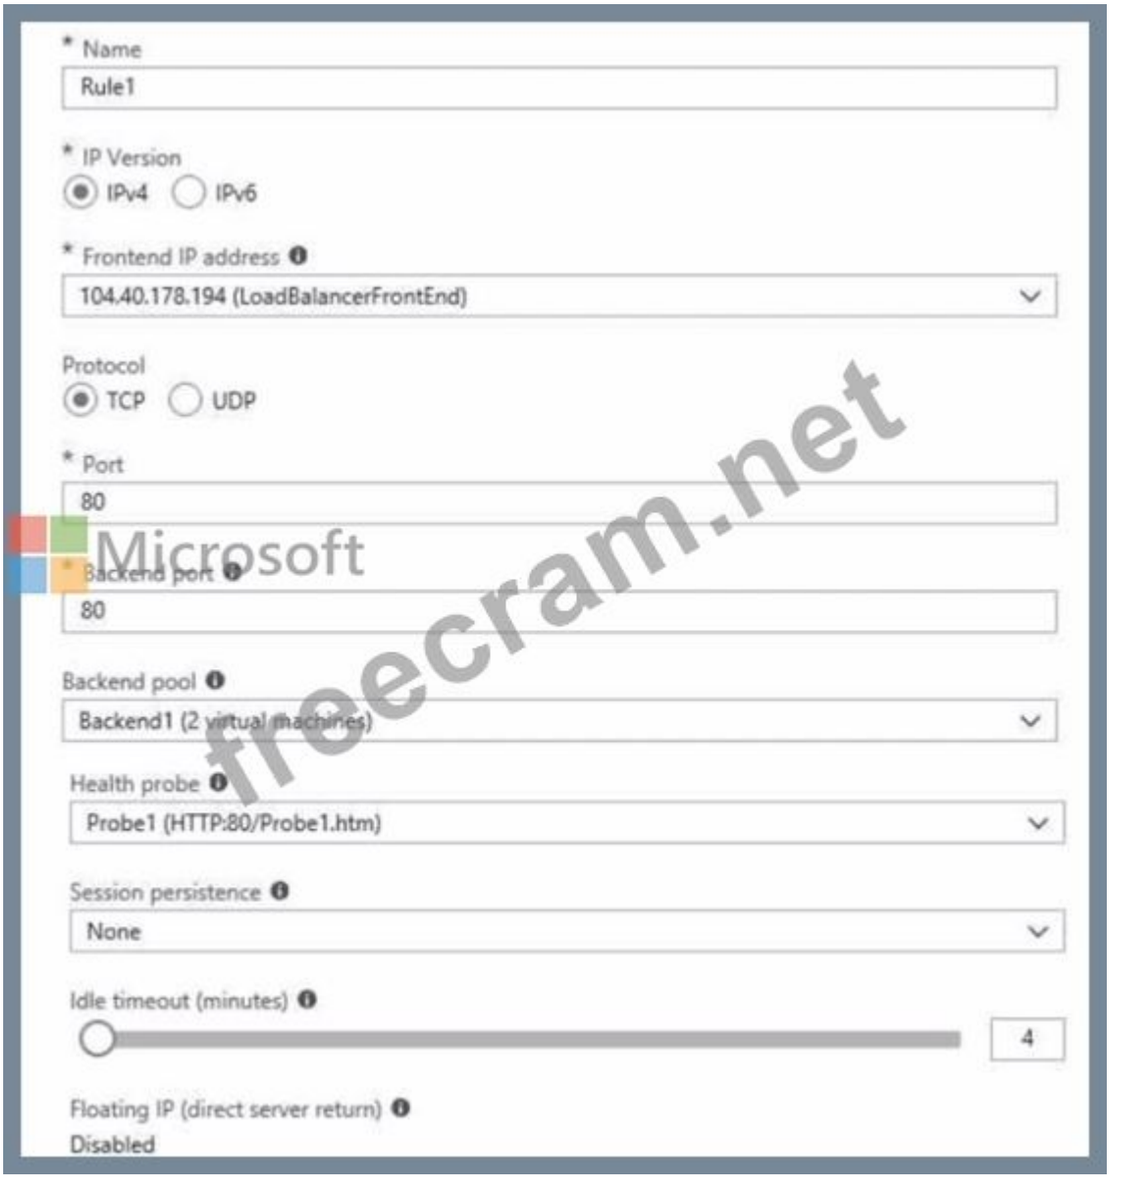

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
VM1 is in the same availability set as VM2.	○	○
If Probe1.htm is present on VM1 and VM2, LB1 will balance traffic for TCP port 80 between VM1 and VM2.	○	○
If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports.	○	○

Answer:

Explanation:

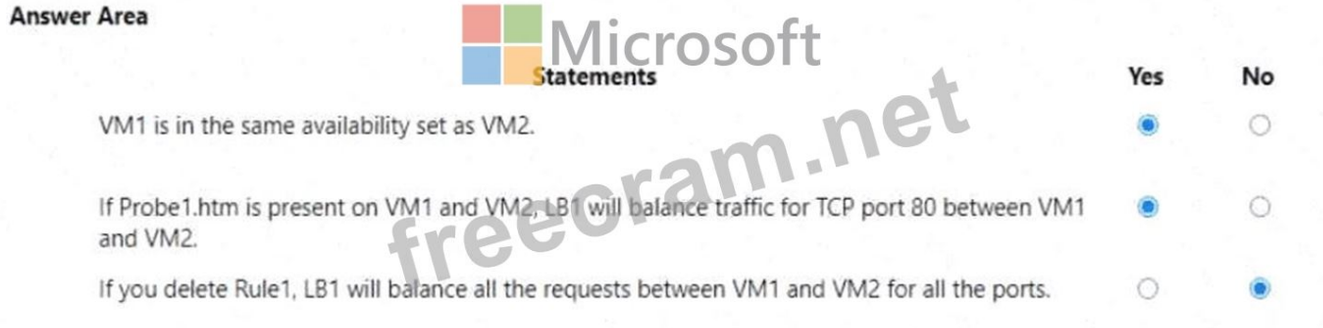

1## VM1 is in the same availability set as VM2 # Yes
In Azure, for a Basic SKU Load Balancer, the backend virtual machines must be in the same availability set (or in the same virtual machine scale set).
Basic Load Balancers cannot span across multiple availability sets or availability zones.
The screenshot shows that LB1 uses the Basic SKU, so by requirement, VM1 and VM2 must be in the same availability set.
Therefore, this statement is true.
2## If Probe1.htm is present on VM1 and VM2, LB1 will balance traffic for TCP port 80 between VM1 and VM2 # Yes A load balancing rule defines how traffic is distributed to the backend pool.
The rule "Rule1" uses TCP port 80 for both frontend and backend.
A health probe (Probe1.htm) monitors the health of backend VMs.
When both VMs respond successfully to Probe1.htm, the load balancer distributes TCP traffic on port 80 evenly between the two healthy VMs.
Therefore, when Probe1.htm is available and healthy on both VMs, LB1 distributes HTTP traffic (port 80) between VM1 and VM2.
This statement is true.
3## If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all ports # No Load balancers in Azure do not automatically balance traffic unless a load-balancing rule exists.
Each load balancing rule explicitly defines which ports and protocols are balanced.
If you delete Rule1, there will be no configuration telling the load balancer which ports to listen to or where to send traffic.
The load balancer will not distribute any traffic automatically for other ports.
Therefore, deleting the rule disables traffic distribution entirely, and this statement is false.

Microsoft Azure Administrator Official Documentation Extract (AZ-104 Reference):
"Basic Load Balancer can only balance traffic among VMs within a single availability set or virtual machine scale set."
"A load-balancing rule defines how traffic is distributed to the backend pool. Deleting a rule stops load balancing for that port."
"Health probes determine which backend pool instances receive new flows." (Source: Microsoft Learn - Azure Load Balancer Overview, Create and Manage Load Balancer Rules, Health Probes in Azure Load Balancer)

NEW QUESTION: 2
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Location
VNET1	Virtual network	East US
IP1	Public IP address	West Europe
RT1	Route table	North Europe

You need to create a network interface named NIC1.
In which location can you create NIC1?
A. East US and North Europe only.
B. East US and West Europe only.
C. East US, West Europe, and North Europe.
D. East US only.
Answer: (SHOW ANSWER)
Before creating a network interface, you must have an existing virtual network in the same location and subscription you create a network interface in.
If you try to create a NIC on a location that does not have any Vnets you will get the following error: "The currently selected subscription and location lack any existing virtual networks. Create a virtual network first." Reference:
https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface

NEW QUESTION: 3
You have an Azure Kubernetes Service (AKS) cluster named AKS1 and a computer named Computer1 that runs Windows 10. Computer1 that has the Azure CLI installed.
You need to install the kubectl client on Computer1.
Which command should you run? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

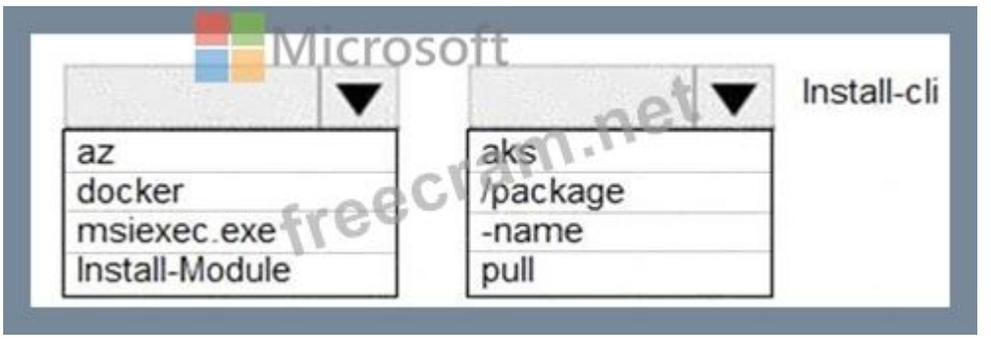

Answer:

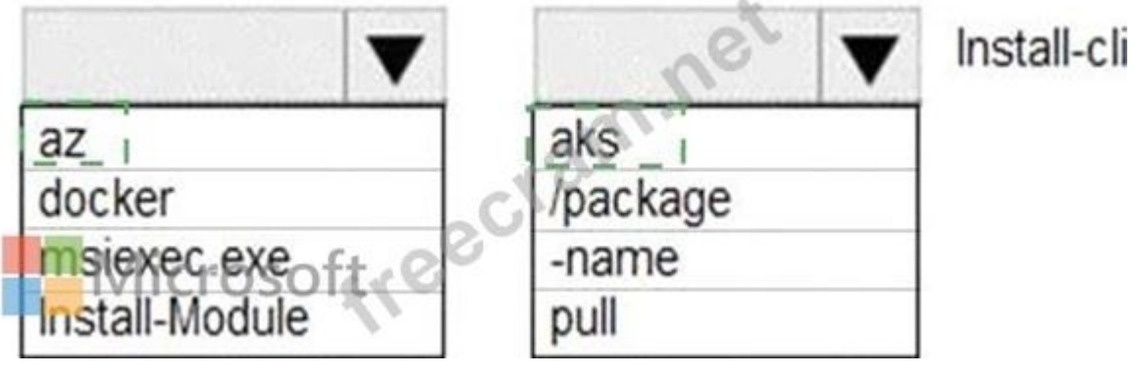

Explanation:

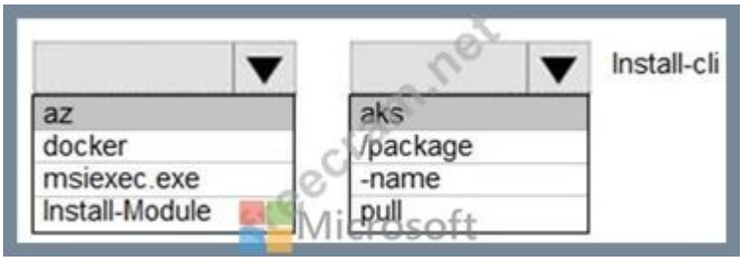

To manage and interact with an Azure Kubernetes Service (AKS) cluster from your local computer, you must have the kubectl client installed. The kubectl command-line tool lets you run commands against Kubernetes clusters to deploy applications, inspect and manage cluster resources, and view logs.

Microsoft provides an official, recommended, and verified way to install kubectl using the Azure CLI.

Correct Command Syntax:

az aks install-cli

Detailed Explanation:

1## Command Context (az)

The az prefix invokes the Azure Command-Line Interface (CLI), which is the unified tool for managing Azure resources directly from a local terminal or script environment.

2## Subcommand (aks)

The aks command group is used to manage Azure Kubernetes Service resources - including creating, configuring, scaling, and connecting to AKS clusters.

3## Action (install-cli)

The install-cli command downloads and installs the latest version of the kubectl binary compatible with your AKS cluster.

When you execute this command, Azure CLI:
Fetches the correct kubectl version that matches your AKS cluster.
Automatically places the binary in a directory included in your system's PATH variable (or instructs you to do so manually).
Ensures version consistency between your management workstation and the AKS cluster.
Alternative (Manual Installation):
While it's possible to install kubectl manually (for example, via Install-Module in PowerShell or using msiexec.exe on Windows), Azure CLI's az aks install-cli method is preferred because it automatically ensures version compatibility with the AKS cluster you are managing.
Microsoft Documentation Extract (Azure Administrator Study Guide Reference):
"To install kubectl using Azure CLI, use the command az aks install-cli. This method ensures the version of kubectl matches the version of Kubernetes used by your AKS cluster." (Source: Microsoft Learn - Manage Azure Kubernetes Service (AKS) with Azure CLI, Module: Deploy and Manage Compute Resources, AZ-104 Exam Guide.)

NEW QUESTION: 4
You have an azure subscription that contains the resources shown in the following table.

Name	Type
VM1	Virtual machine
VNet1	Virtual network
NIC1	Network interface
LB1	Load balancer
VPN1	Virtual network gateway

You create a public IP address named IPI.
Which two resources can you associate to IP1.
Each correct answer presents a complete solution
NOTE: Each correct selection is worth one point
A. VM1
B. NIC1
C. VPN1
D. LB1
E. VNet1
Answer: (SHOW ANSWER)
B). NIC1 (Network Interface): Public IP addresses are directly associated with network interfaces. This allows the virtual machine connected to that interface to communicate with the internet.
D). LB1 (Load Balancer): Load balancers can have a public IP address assigned to their frontend configuration. This allows external traffic to reach the load balancer, which then distributes it to the backend VMs.

NEW QUESTION: 5
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, those questions will not appear in the review screen.
You have a Microsoft Entra tenant named contoso.com.
You have a CSV file that contains the names and email addresses of 500 external users.
You need to create a guest user account in contoso.com for each of the 500 external users.
Solution; From Microsoft Entra ID in the Azure portal, you use the Bulk create user operation.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In Microsoft Entra ID (Azure Active Directory), creating guest user accounts for external users involves the B2B collaboration invitation process, not the Bulk create user operation. The Bulk create user feature in the Azure portal is used exclusively for creating member users (internal accounts) within your tenant - users that belong directly to your organization.
The Microsoft Azure Administrator official study guide and documentation specify that to onboard external users (guests), you must invite them using the B2B invitation mechanism. This process creates user objects with userType = Guest in your directory. The correct method to achieve this in bulk is by using the PowerShell cmdlet New-AzureADMSInvitation or Microsoft Graph API to send invitations automatically to all external users listed in your CSV file.
Bulk create user operation does not support adding guest accounts because it doesn't generate the B2B invitation link or email, which is required for the external user to accept and authenticate using their own identity provider (e.g., Microsoft account, Google, or another Entra tenant).
Therefore, the proposed solution fails to meet the requirement of creating guest accounts for the 500 external users in the contoso.com tenant. The correct solution would be to use PowerShell with the New- AzureADMSInvitation command or the Azure AD portal's "Invite external users" bulk import option designed specifically for guest creation.

NEW QUESTION: 6
You have an Azure subscription named Sub1.
You plan to deploy a multi-tiered application that will contain the tiers shown in the following table.

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

You need to recommend a networking solution to meet the following requirements:
* Ensure that communication between the web servers and the business logic tier spreads equally across the virtual machines.

* Protect the web servers from SQL injection attacks.

Which Azure resource should you recommend for each requirement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

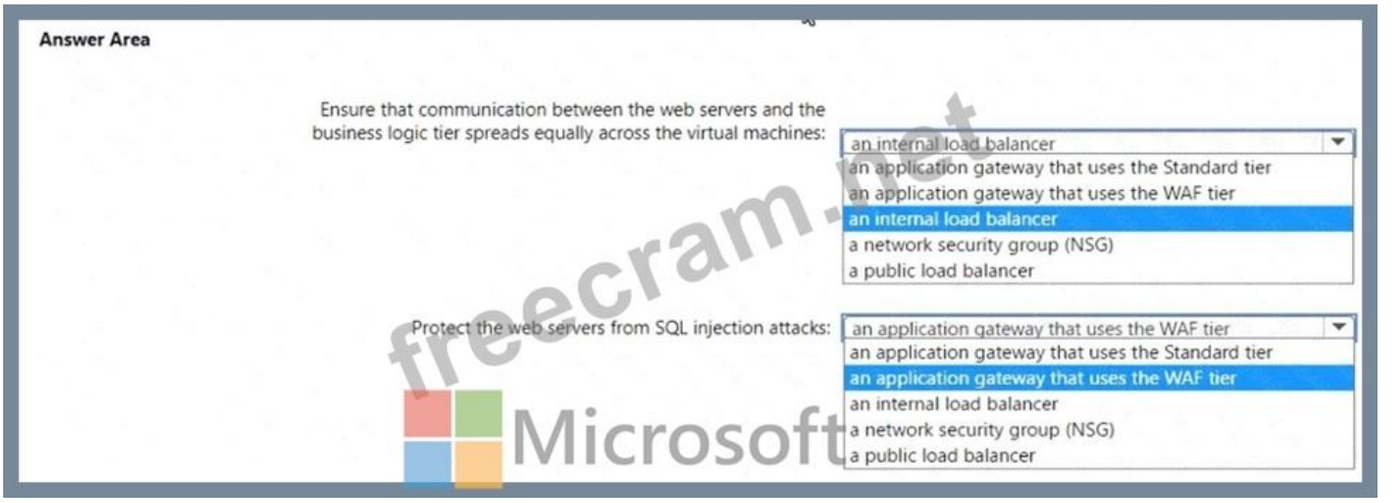

Answer:

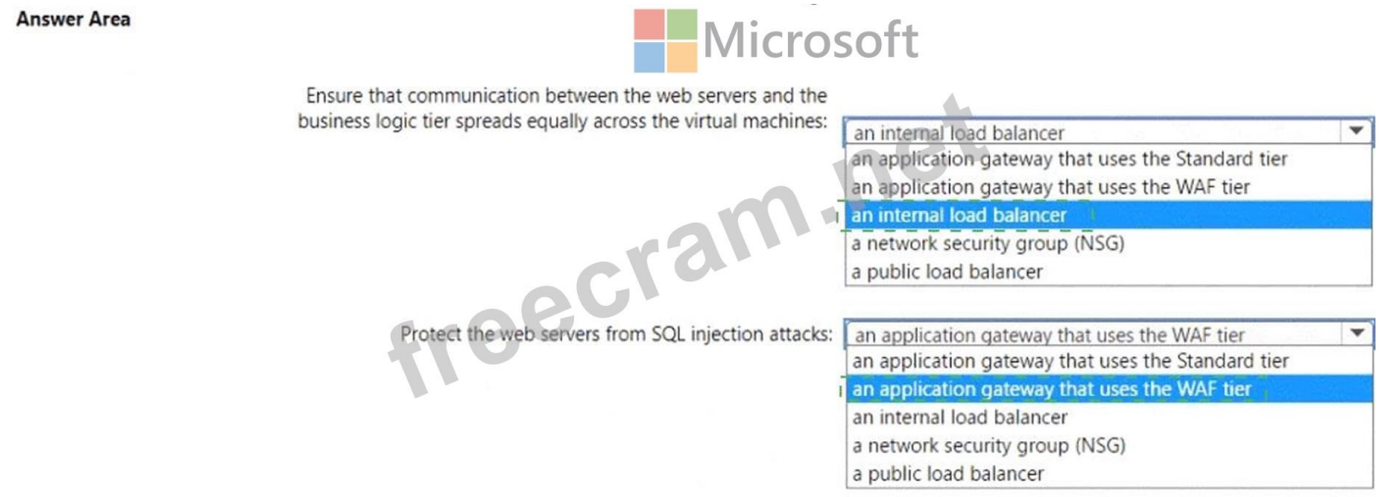

Explanation:

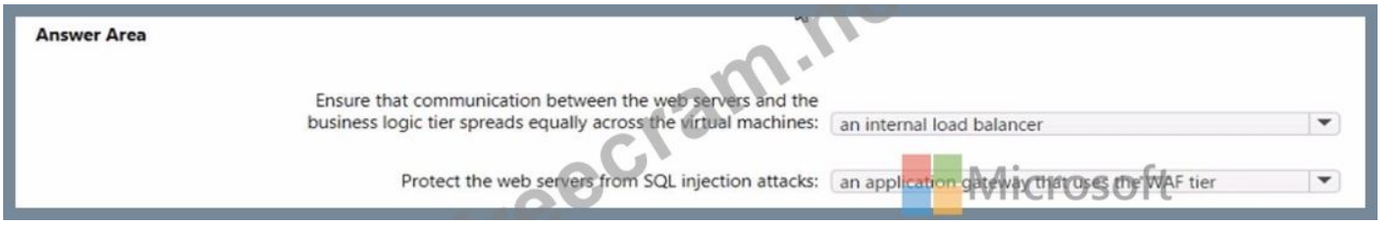

Box 1: an internal load balancer

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: 7

You have an Azure subscription named Sub1 that contains the resources shown in the following table.

Name	Description
RG1	Resource group
VNet1	Virtual network in RG1

The subscription contains the users shown in the following table.

Name	Role	Role scope
Admin1	Owner	Sub1
Admin2	Contributor	RG1

You have the following Bicep file named Deploy.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'
        }
      }
    ]
  }
}
```


You run the following command.

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

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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

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

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: 8

You need to identify which storage account to use for the flow logging of IP traffic from VM5. The solution must meet the retention requirements.

Which storage account should you identify?

A. storage4

B. storage1

C. storage2

D. storage3

Answer: C (LEAVE A REPLY)

You must identify the correct storage account for flow logging of IP traffic from VM5 that satisfies the retention requirement of eight months.

Step 1 - Review the Requirements

Flow logs are stored in a StorageV2 (general purpose v2) account.

The selected storage account must support Network Watcher flow logs.

Data must be retained for eight months (# 240 days).

Step 2 - Evaluate Each Storage Account

Storage Account

Kind

Region

Remarks

storage1

Storage (general purpose v1)

West US

Does not support flow logs (deprecated type).

storage2

StorageV2 (general purpose v2)

East US

Fully supports flow logs and lifecycle management.

storage3

BlobStorage

East US 2

Not suitable - supports blobs only, not log structure or lifecycle retention.

storage4

FileStorage

Central US

File-only - cannot store flow logs.

Step 3 - Compliance with Flow Log Retention

Flow logs for Network Watcher NSG flow logging are supported only by StorageV2 accounts. You can use Azure Storage lifecycle management to automatically delete logs after a set period - such as 240 days (8 months) - to comply with retention requirements.

Official Microsoft Note:

"Network security group (NSG) flow logs are stored in Azure StorageV2 (General Purpose v2) accounts, which support lifecycle management for log retention." - Azure Network Watcher documentation.

Conclusion
storage2 is the only StorageV2 account.
It's located in East US, matching VM5's region (East US) - this satisfies performance and compliance requirements.

NEW QUESTION: 9
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You need to ensure that an Azure Active Directory (Azure AD) user named Admin1 is assigned the required role to enable Traffic Analytics for an Azure subscription.
Solution: You assign the Traffic Manager Contributor role at the subscription level to Admin1
A. Yes
B. NO
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: 10
You have an Azure subscription that contains the virtual machines shown in the following table.

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

VM1 and VM2 use public IP addresses. From Windows Server 2019 on VM1 and VM2, you allow inbound Remote Desktop connections.
Subnet1 and Subnet2 are in a virtual network named VNET1.
The subscription contains two network security groups (NSGs) named NSG1 and NSG2. NSG1 uses only the default rules.
NSG2 uses the default rules and the following custom incoming rule;
* Priority: 100
* Name: Rule1
* Port: 3389
* Protocol: TCP
* Source: Any
* Destination: Any
* Action: Allow
NSG1 is associated to Subnet! NSG2 is associated to the network interface of VM2.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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:

Answer Area

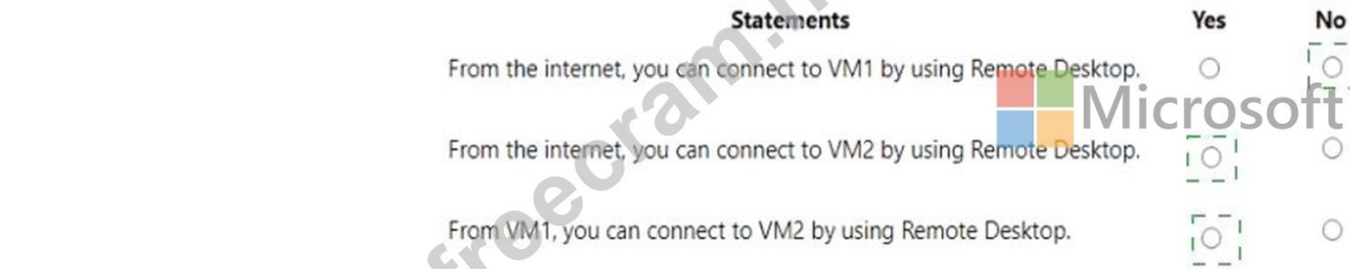

Explanation:

Answer Area

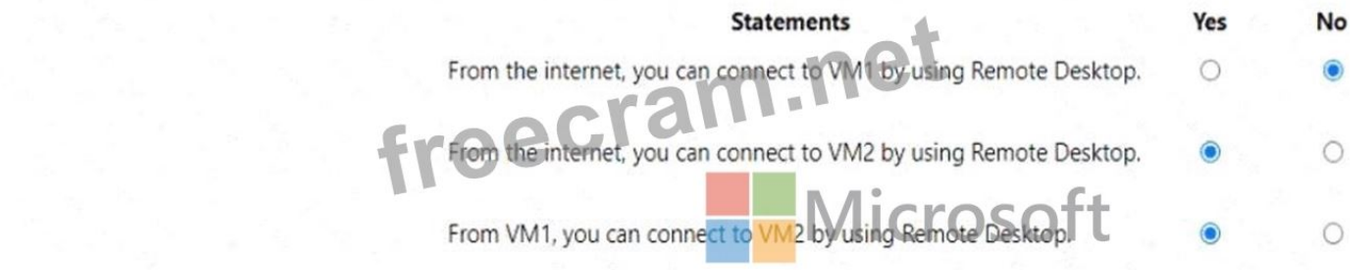

In Azure, Network Security Groups (NSGs) control inbound and outbound traffic to network interfaces (NICs), subnets, and virtual machines (VMs) using rules based on priority and direction.

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: 11
You have an Azure subscription named Subscription1 that contains the resources shown in the following table.

Name	Type	Region	Resource group
RG1	Resource group	West Europe	*Not applicable*
RG2	Resource group	North Europe	*Not applicable*
Vault1	Recovery Services vault	West Europe	RG1

You create virtual machines in Subscription1 as shown in the following table.

Name	Resource group	Region	Operating system
VM1	RG1	West Europe	Windows Server 2016
VM2	RG1	North Europe	Windows Server 2016
VM3	RG2	West Europe	Windows Server 2016
VMA	RG1	West Europe	Ubuntu Server 18.04
VMB	RG1	North Europe	Ubuntu Server 18.04
VMC	RG2	West Europe	Ubuntu Server 18.04

You plan to use Vault1 for the backup of as many virtual machines as possible.
Which virtual machines can be backed up to Vault1?
A. VM1, VM3, VMA, and VMC only
B. VM1 and VM3 only
C. VM1, VM2, VM3, VMA, VMB, and VMC
D. VM1 only
E. VM3 and VMC only
Answer: A (LEAVE A REPLY)
To create a vault to protect virtual machines, the vault must be in the same region as the virtual machines. If you have virtual machines in several regions, create a Recovery Services vault in each region.
References:
https://docs.microsoft.com/bs-cyrl-ba/azure/backup/backup-create-rs-vault

NEW QUESTION: 12
You have an Azure subscription.
You plan to migrate 50 virtual machines from VMware vSphere to the subscription.

You create a Recovery Services vault.
What should you do next?
A. Configure an extended network.
B. Create a recovery plan.
C. Deploy an Open Virtualization Application (OVA) template to vSphere.
D. Configure a virtual network.
Answer: (SHOW ANSWER)
To migrate virtual machines from VMware vSphere to Azure, you need to use Azure Migrate, which is a service that helps you assess and migrate your on-premises workloads to Azure. Azure Migrate uses an appliance that you deploy as an Open Virtualization Application (OVA) template to vSphere. The appliance discovers the virtual machines and sends metadata and performance data to Azure Migrate. You can then use Azure Migrate to assess the readiness, cost, and sizing of the virtual machines for migration. You can also use Azure Migrate to replicate and migrate the virtual machines to Azure. References:
* About Azure Migrate
* Prepare VMware servers for assessment and migration to Azure with Azure Migrate Server Migration

NEW QUESTION: 13
You have an Azure subscription that contains the public load balancers shown in the following table.

Name	SKU
LB1	Basic
LB2	Standard

You plan to create six virtual machines and to load balance requests to the virtual machines. Each load balancer will load balance three virtual machines.
You need to create the virtual machines for the planned solution.
How should you create the virtual machines? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

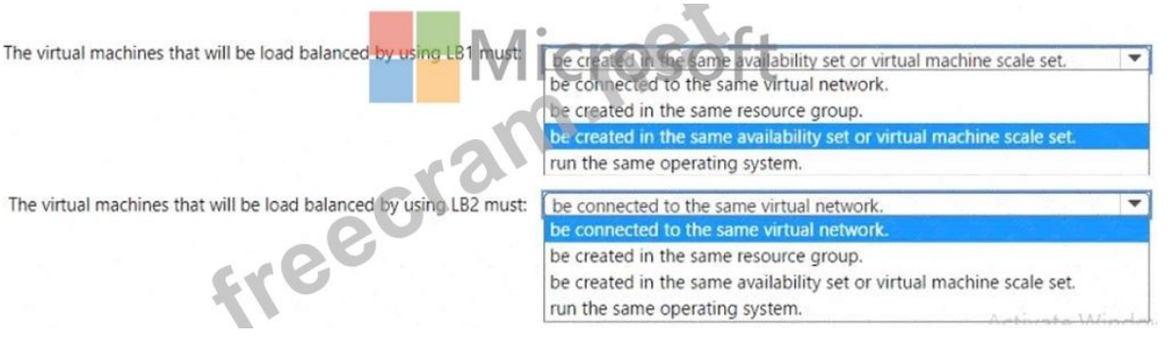

Answer:

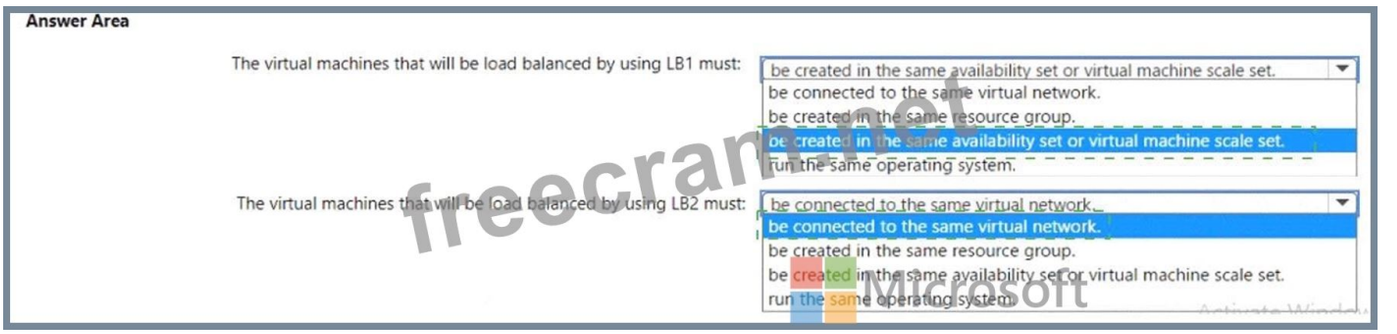

Explanation:

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: 14
You have an Azure subscription that contains a storage account named storage 1 in the North Europe A2ure region.
You need to ensure that when blob data is added to storage1, a secondary copy is created in the East US region. The solution must minimize administrative effort.
What should you configure?
A. operational backup
B. a lifecycle management rule
C. object replication
D. geo-redundant storage (GRS)
Answer: (SHOW ANSWER)
In Microsoft Azure Storage, maintaining high availability and durability of data is achieved through replication. There are multiple redundancy options - LRS (Locally Redundant Storage), ZRS (Zone- Redundant Storage), GRS (Geo-Redundant Storage), and RA-GRS (Read-Access Geo-Redundant Storage) - each serving different scenarios.
However, when the requirement is to create a secondary copy of blob data in a specific region (for example, when blob data is added in North Europe, a copy must be created in East US), geo-redundant storage (GRS) cannot be used directly because GRS replicates data automatically to a paired region, not to a region of your choice. In this case, North Europe's paired region is West Europe, not East US. Therefore, GRS does not meet the goal.
The Microsoft Azure Administrator documentation specifies that the correct method for replicating data between user-defined regions is Object Replication for Block Blobs (OR).
Object replication allows you to:
* Automatically replicate new and updated block blobs between two storage accounts in different Azure regions.
* Use replication policies to define source and destination containers.
* Minimize administrative effort, as once configured, replication occurs automatically and asynchronously.
* Retain independent storage accounts, ensuring flexibility and control over access and lifecycle management.
Unlike lifecycle management rules (which only move or delete blobs within the same storage account) or operational backup (which protects data for restore purposes, not replication), object replication precisely meets the requirement of maintaining a secondary copy in another chosen region.

Therefore, per Azure's official Storage replication and Object replication documentation, the correct configuration is Object replication between storage1 (North Europe) and another storage account in East US.

NEW QUESTION: 15
You create a Recovery Services vault backup policy named Policy1 as shown in the following exhibit.

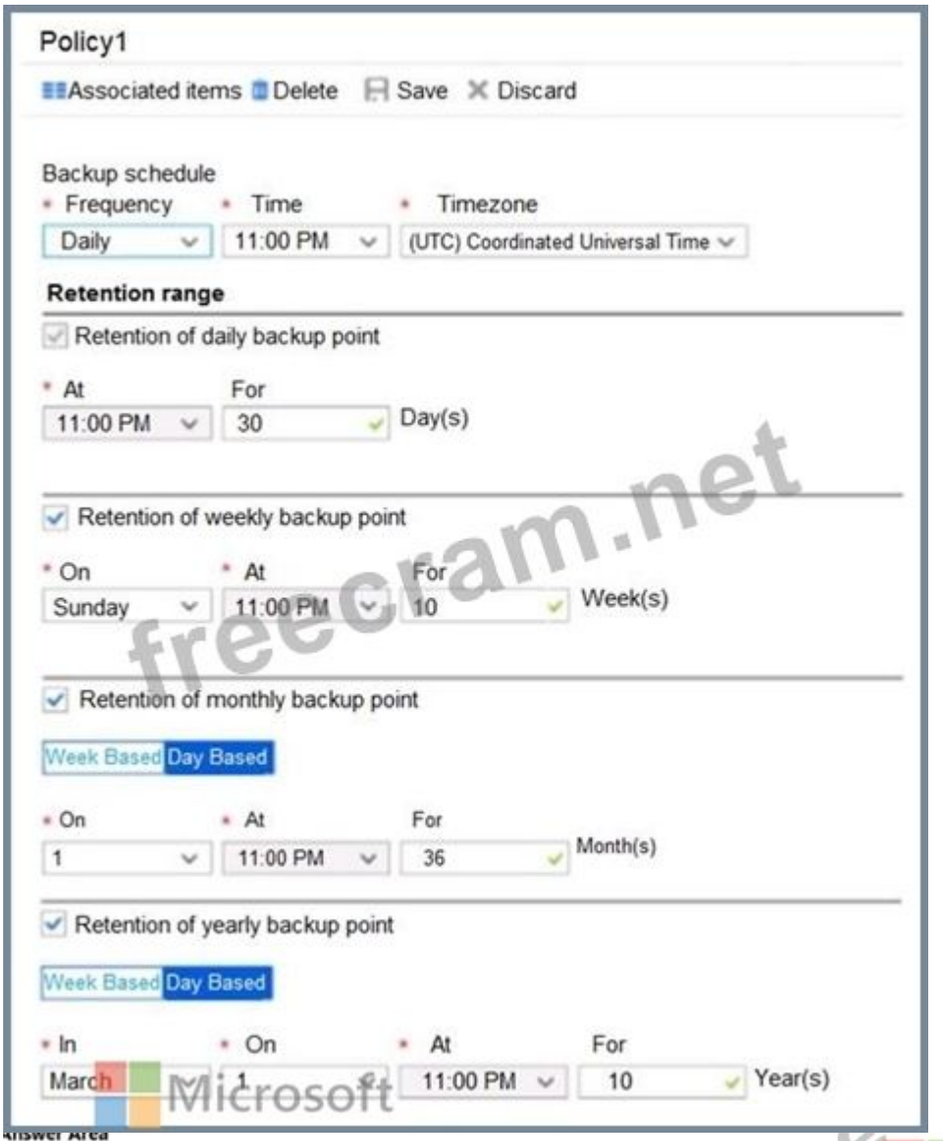

Answer Area

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

30 days
10 weeks
36 months
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:

Answer Area

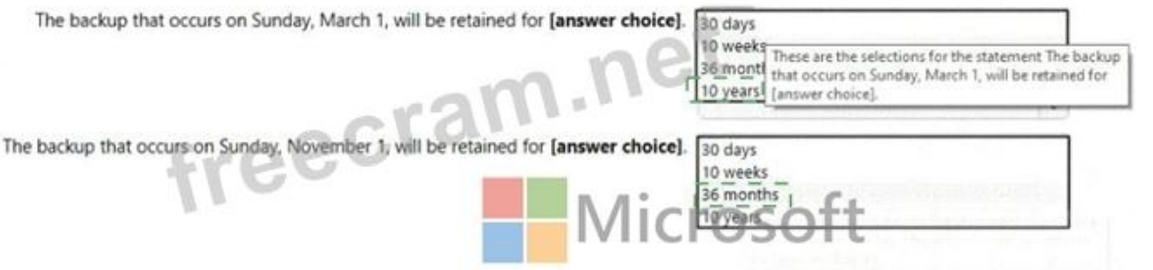

Explanation:
Box 1: 10 years
The yearly backup point occurs to 1 March and its retention period is 10 years.
Box 2: 36 months
The monthly backup point occurs on the 1
of every month and its retention period is 36 months.
Note: Azure retention policy takes the longest period of retention for each backup. In case of conflict between
2 different policies.
Reference:
https://docs.microsoft.com/en-us/microsoft-365/compliance/retention?view=o365-worldwide

NEW QUESTION: 16
You have an on-premises network.
You have an Azure subscription that contains three virtual networks named VNET1, VNET2, and VNET3.
The virtual networks are peered and connected to the on-premises network. The subscription contains the virtual machines shown in the following table.

Name	Region	Connected to
VM1	West US	VNET1
VM2	West US	VNET1
VM3	West US	VNET2
VM4	Central US	VNET3

You need to monitor connectivity between the virtual machines and the on-premises network by using Connection Monitor. What is the minimum number of connection monitors you should deploy?
A. 1
B. 2
C. 3
D. 4
Answer: (SHOW ANSWER)
Azure Connection Monitor (a feature within Network Watcher) provides a unified, end-to-end monitoring experience across Azure, on-premises, and hybrid environments. It enables administrators to monitor connectivity between virtual machines, endpoints, and on-premises

networks without needing separate monitors per region when all resources can be accessed through peering or connected networks.
According to the Microsoft Azure Network Watcher documentation,
"Connection Monitor enables you to monitor network communication between a virtual machine (VM) or virtual machine scale set and any endpoint. These endpoints can be in Azure, on-premises, or any external environment. A single Connection Monitor can track connectivity across peered VNets and hybrid connections." In this case:
* VNET1, VNET2, and VNET3 are all peered and connected to the on-premises network.
* Peered VNets provide full IP-level connectivity between all VMs as if they were on the same network, subject to NSG and routing rules.
* The on-premises network connection allows visibility for hybrid monitoring.
Therefore, from any region or network where Network Watcher is enabled, you can create a single Connection Monitor that includes:
* Source endpoints (VM1, VM2, VM3, and VM4) from any VNet (since they are peered and reachable).
* Destination endpoint (the on-premises network or any IP/FQDN).
Connection Monitor can use Test Groups to define multiple source-destination pairs within a single monitor.
This eliminates the need to deploy multiple Connection Monitors per region as long as the sources are discoverable through VNet peering or hybrid connections.
Microsoft Official Documentation Extract (Azure Network Watcher - Connection Monitor Overview):
"Connection Monitor provides unified end-to-end monitoring between Azure and on-premises resources. It supports monitoring across virtual networks, peered virtual networks, and hybrid environments using a single monitor instance." (Microsoft Learn: Azure Network Watcher - Connection Monitor Overview, "Unified Connection Monitoring" section)

NEW QUESTION: 17
You have an Azure subscription that contains the virtual networks shown in the following table.

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

Each virtual network has 50 connected virtual machines.

You need to implement Azure Bastion. The solution must meet the following requirements:

* Support host scaling.
* Support uploading and downloading files.
* Support the virtual machines on both VNet1 and VNet2.
* Minimize the number of addresses on the Azure Bastion subnet.

How should you configure Azure Bastion? To answer, select the options in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

Answer:

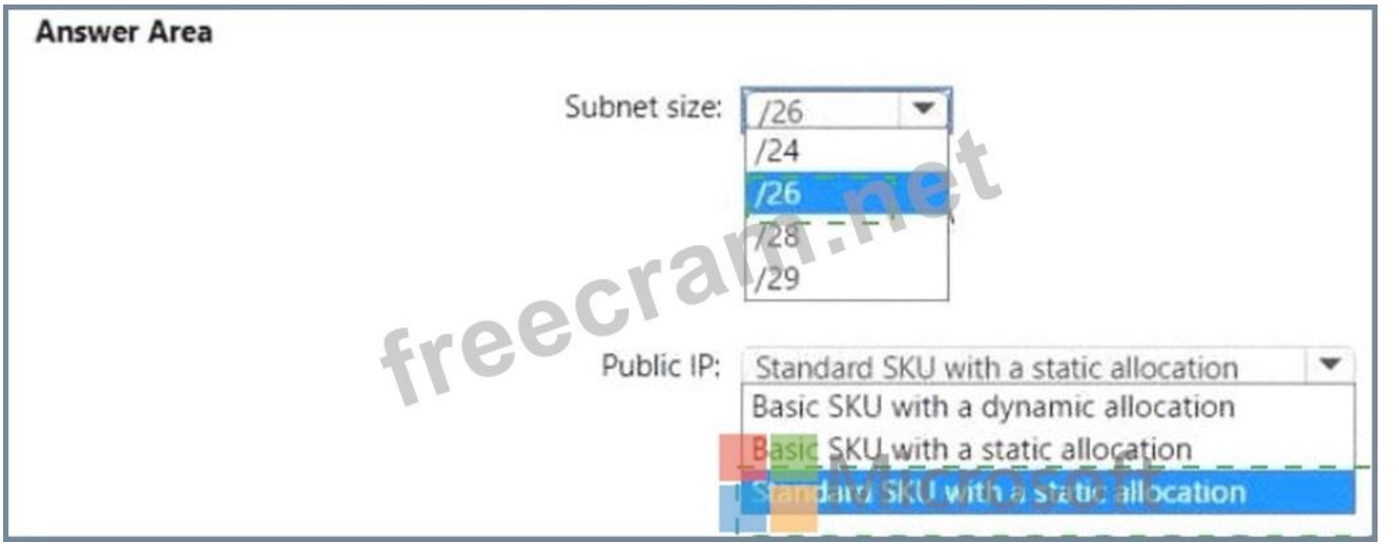

Explanation:

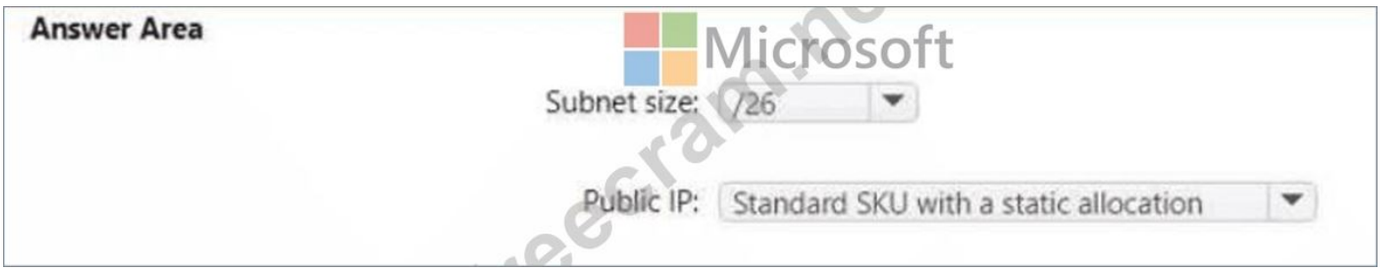

Azure Bastion is a fully managed service that provides secure and seamless RDP/SSH connectivity to your virtual machines directly through the Azure portal - without exposing those VMs to public IP addresses.

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: 18
You have an Azure subscription named Sub1 that contains two users named User1 and User2.
You need to assign role-based access control (RBAC) roles to User1 and User2. The users must be able to perform the following tasks in Sub1:
* User1 must view the data in any storage account.
* User2 must assign users the Contributor role for storage accounts.
The solution must use the principle of least privilege.
Which RBAC role should you assign to each user? To answer, drag the appropriate roles to the correct users.
Each role may be used once, more than once, or not at all.

Answer:

Explanation:
User1: You should assign the Reader and Data Access role to User1. This role grants read access to Azure resources and data, including the data in any storage account1. This role is suitable for User1's task of viewing the data in any storage account, and it follows the principle of least privilege by not granting any write or delete permissions.
User2: You should assign the Storage Account Contributor role to User2. This role grants full access to manage storage accounts and their data, including the ability to assign roles in Azure RBAC2. This role is suitable for User2's task of assigning users the Contributor role for storage accounts, and it follows the principle of least privilege by not granting access to other types of resources.

NEW QUESTION: 19
You create an Azure VM named VM1 that runs Windows Server 2019.
VM1 is configured as shown in the exhibit (Click the Exhibit tab.)

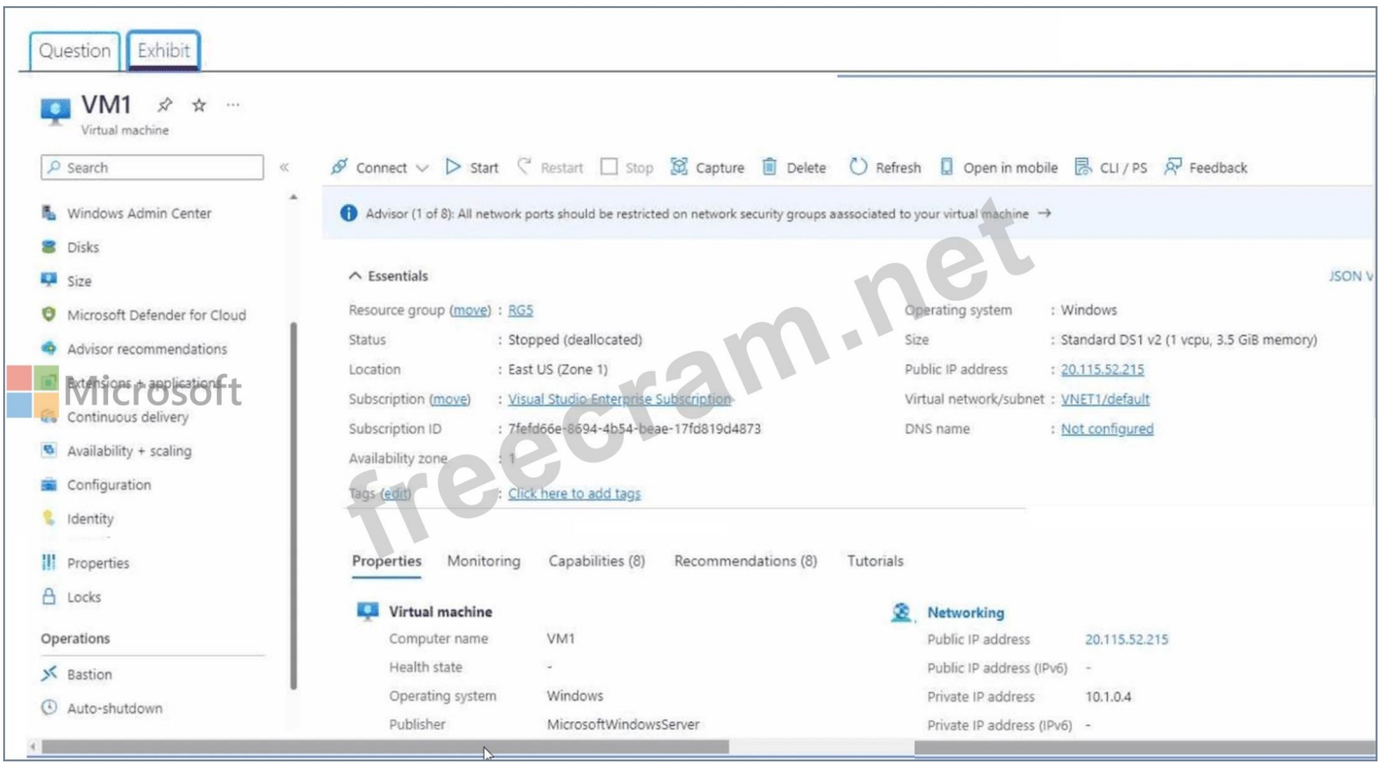

You need to enable Desired State Configuration for VM1.

What should you do first?

A. Configure a DNS name for VM1.

B. Start VM1.

C. Capture a snapshot of VM1.

D. Connect to VM1.

Answer: B (LEAVE A REPLY)

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: 20
You have two Azure virtual machines as shown in the following table.

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

You create the Azure DNS zones shown in the following table.

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

You perform the following actions:
To fabrikam.com, you add a virtual network link to vnet1 and enable auto registration.
For contoso.com, you assign vm1 and vm2 the Owner role.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worm one point.

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: 21

You have 15 Azure subscriptions.

You have a Microsoft Entra tenant that contains a security group named Group1.

You plan to purchase additional Azure subscriptions.

You need to ensure that Group1 can manage role assignments for the existing subscriptions and the planned subscriptions. The solution must meet the following requirements:

* Use the principle of least privilege.

* Minimize administrative effort

What should you do?

A. Create a new management group and assign Group1 the User Access Administrator role for the group.

B. Assign Group1 the Owner role for the root management group.

C. Assign Group1 the User Access Administrator role for the root management group.

D. Create a new management group and assign Group1 the Owner role for the group.

Answer: (SHOW ANSWER)

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: 22
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
VNET1	Virtual network	Contains subnet1 and subnet2
subnet1	Subnet	IP address space 10.3.0.0/24
subnet2	Subnet	IP address space 10.4.0.0/24
NSG1	Network security group (NSG)	*None*
vm1	Virtual machine	IP address 10.3.0.15
vm2	Virtual machine	IP address 10.4.0.16
storage1	Storage account	*None*

NSG1 is configured as shown in the following exhibit.

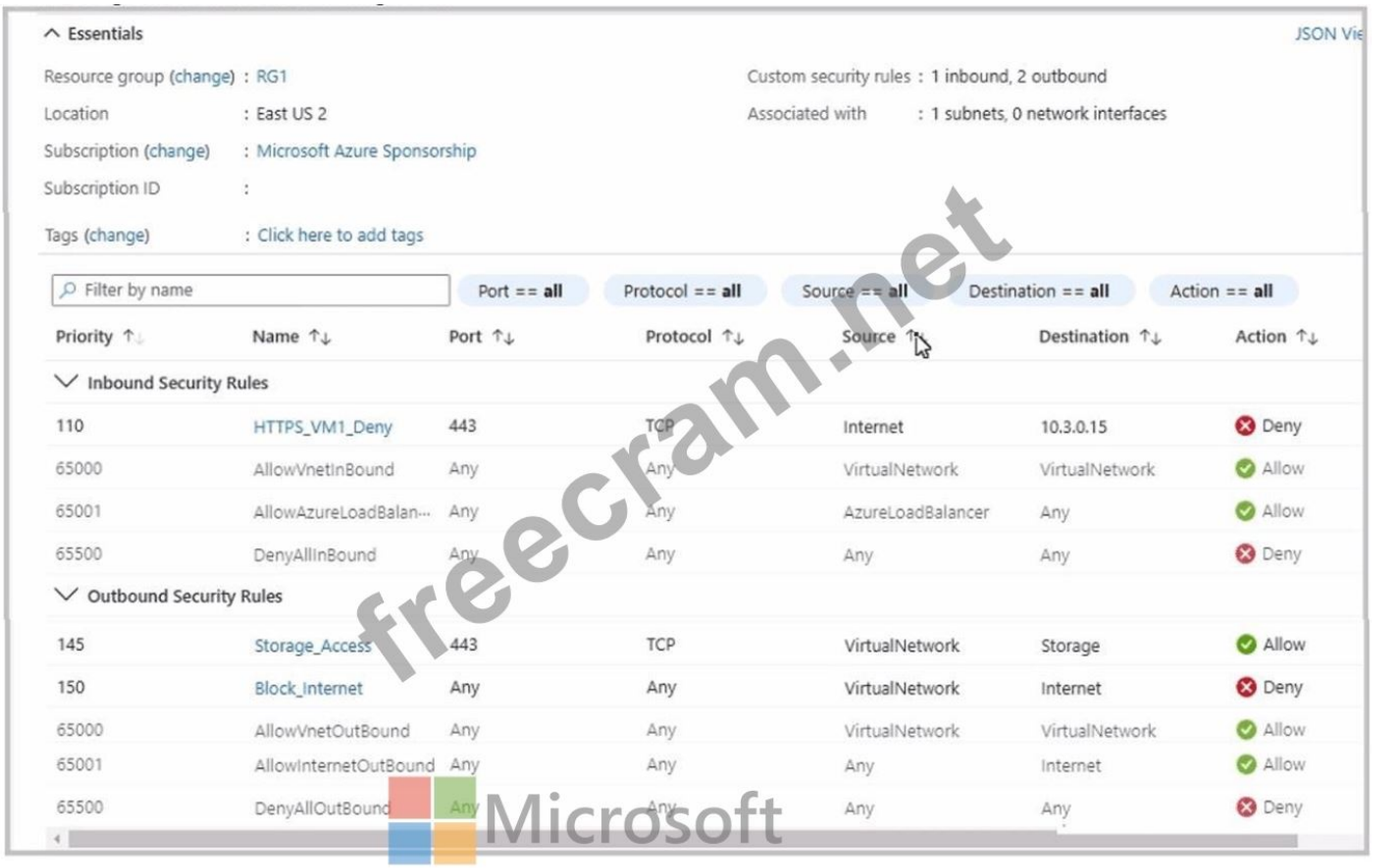

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
VM1 can access storage1.	○	○
VM2 can access VM1 by using the HTTPS protocol.	○	○
The security rules for NSG1 apply to any virtual machine on VNET1.	○	○

Answer:

Answer Area

Statements	Yes	No
VM1 can access storage1.	○	○
VM2 can access VM1 by using the HTTPS protocol.	○	○
The security rules for NSG1 apply to any virtual machine on VNET1.	○	○

Explanation:

Yes - VM1 can access the Storage account because there is nothing blocking it the on the virtual network.

There is a rule that actually allows outbound access to storage.

Yes- VM2 is on the Same VNET there is nothing blocking access to it from VM1 on the Virtual network. The Deny rule for HTTPS_VM1_Deny is for inbound connections from the internet.
No- You have a Inbound deny rule for VM1 from the the internet with a destination of the 10.3.0.15 which is in Subnet1. This proves the NSG is associated to Subnet1 and only subnet one because the image shows it is connected to only 1 subnet. VM2 is on Subnet2 which you can determined by its IP address. This means that NSG1 does not apply to VM2.

NEW QUESTION: 23
You have an Azure subscription named Subscription1.
In Subscription1, you create an Azure file share named share1.
You create a shared access signature (SAS) named SAS1 as shown in the following exhibit.

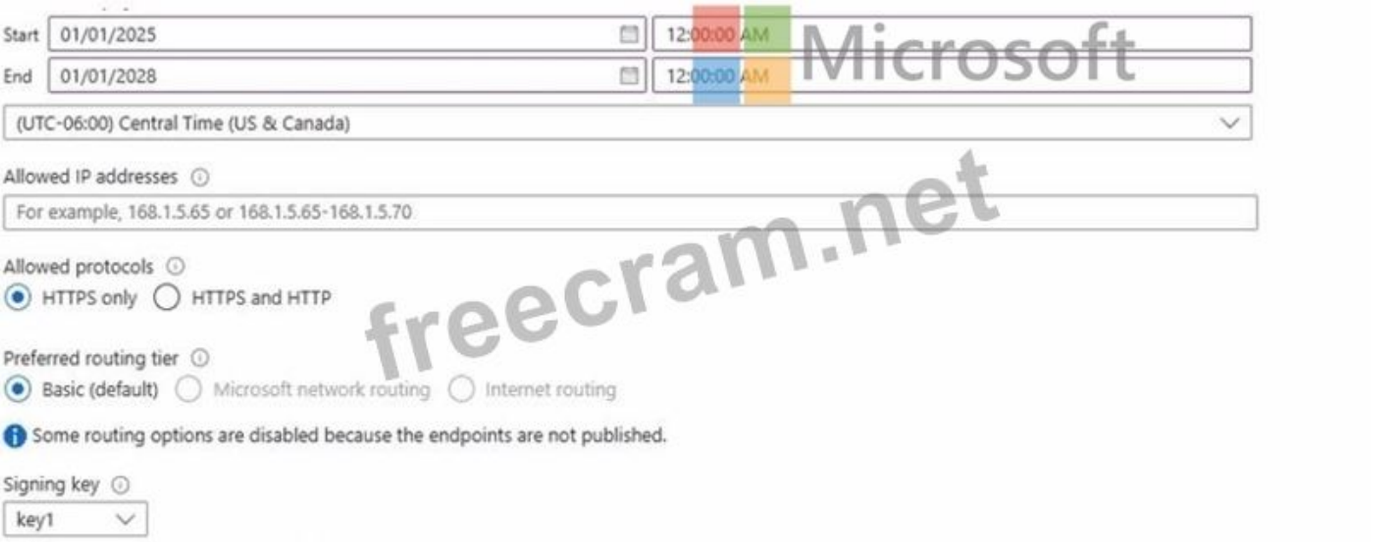

To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

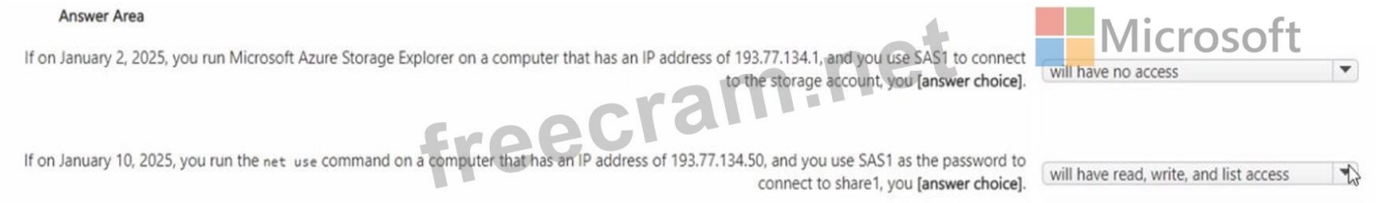

A Shared Access Signature (SAS) defines how, when, and from where a storage resource can be accessed.

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: 24
You have an Azure subscription. The subscription contains virtual machines that run Windows Server.
You have a data collection rule (DCR) named Rule1
You plan to use the Azure Monitor Agent to collect events from Windows System event logs.
You only need to collect system events that have an ID of 1001.
Which type of query should you use for the data source in Rulel?
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: 25
You have an Azure subscription that contains a container group named Group1. Group1 contains two Azure container instances as shown in the following table.

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

You need to ensure that container2 can use CPU resources without negatively affecting container1.
What should you do?
A. Decrease the resource limit of container2 to two CPUs.
B. Remove the resource limit for both containers.
C. Increase the resource limit of container1 to three CPUs.
D. Increase the resource limit of container2 to six CPUs.
Answer: (SHOW ANSWER)

NEW QUESTION: 26
You have an Azure subscription. The subscription contains a virtual machine that runs Windows 10.
You need to join the virtual machine to an Active Directory domain.

How should you complete the Azure Resource Manager (ARM) template? To answer, select the appropriate options in the answer area.

NOTE Each correct selection is worth one point.

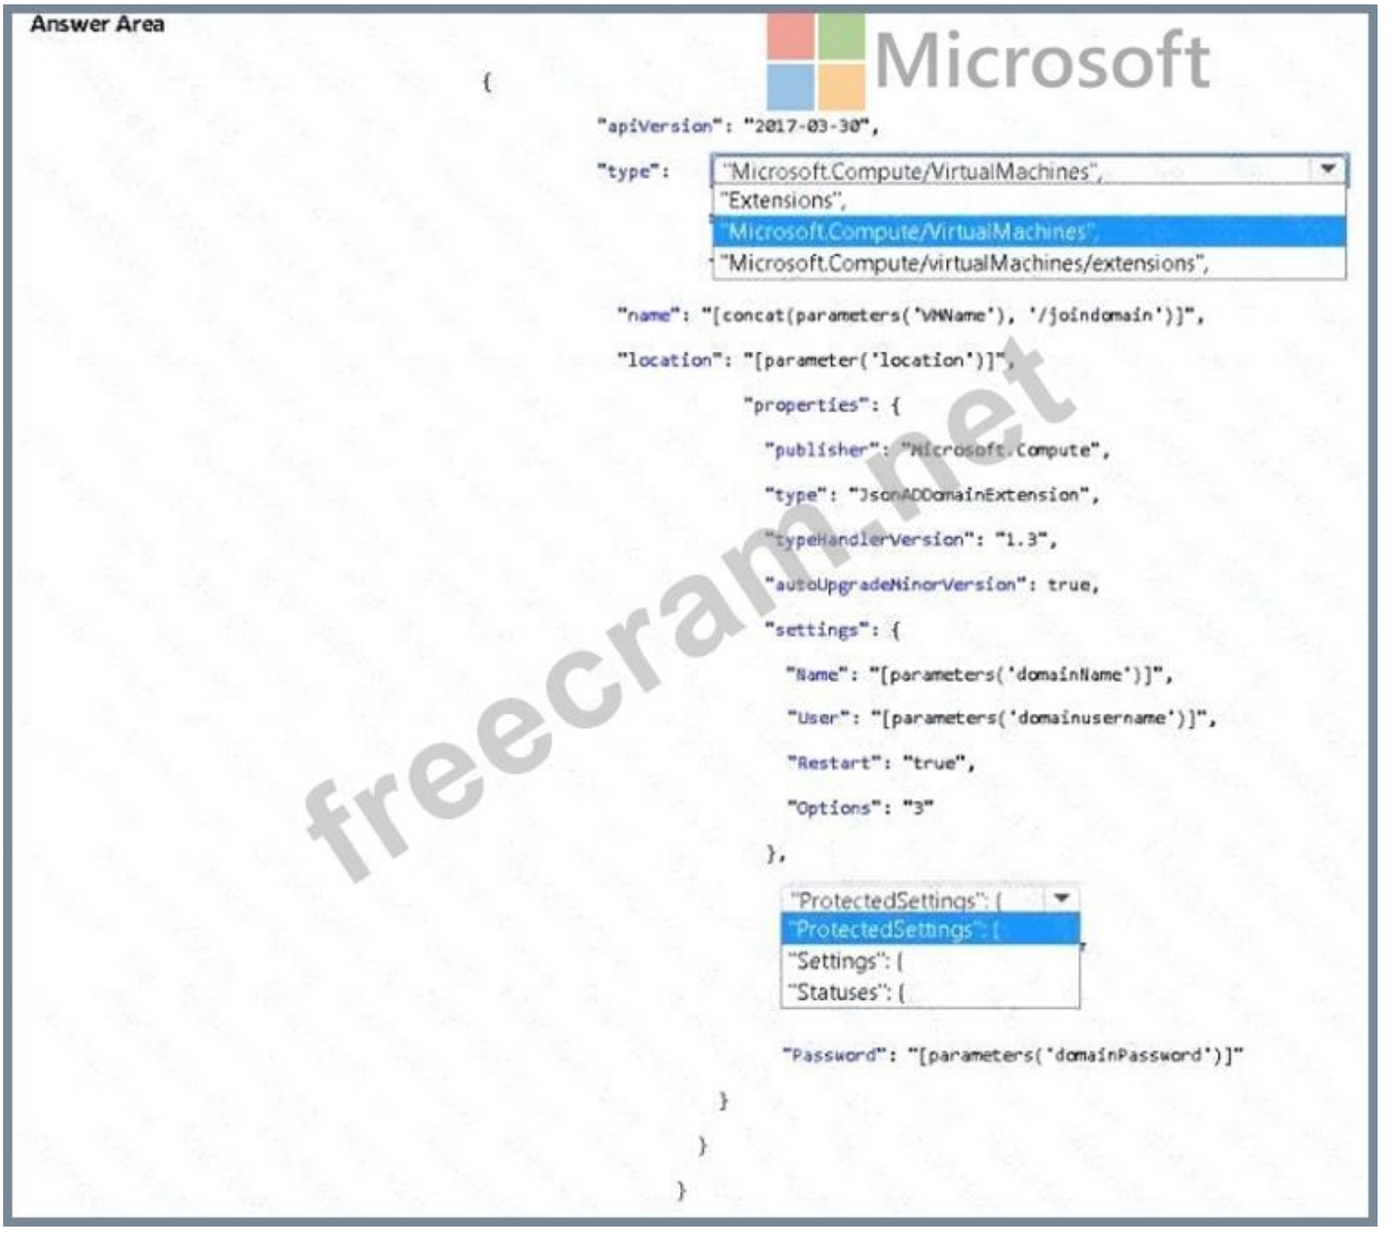

Answer:

Answer Area

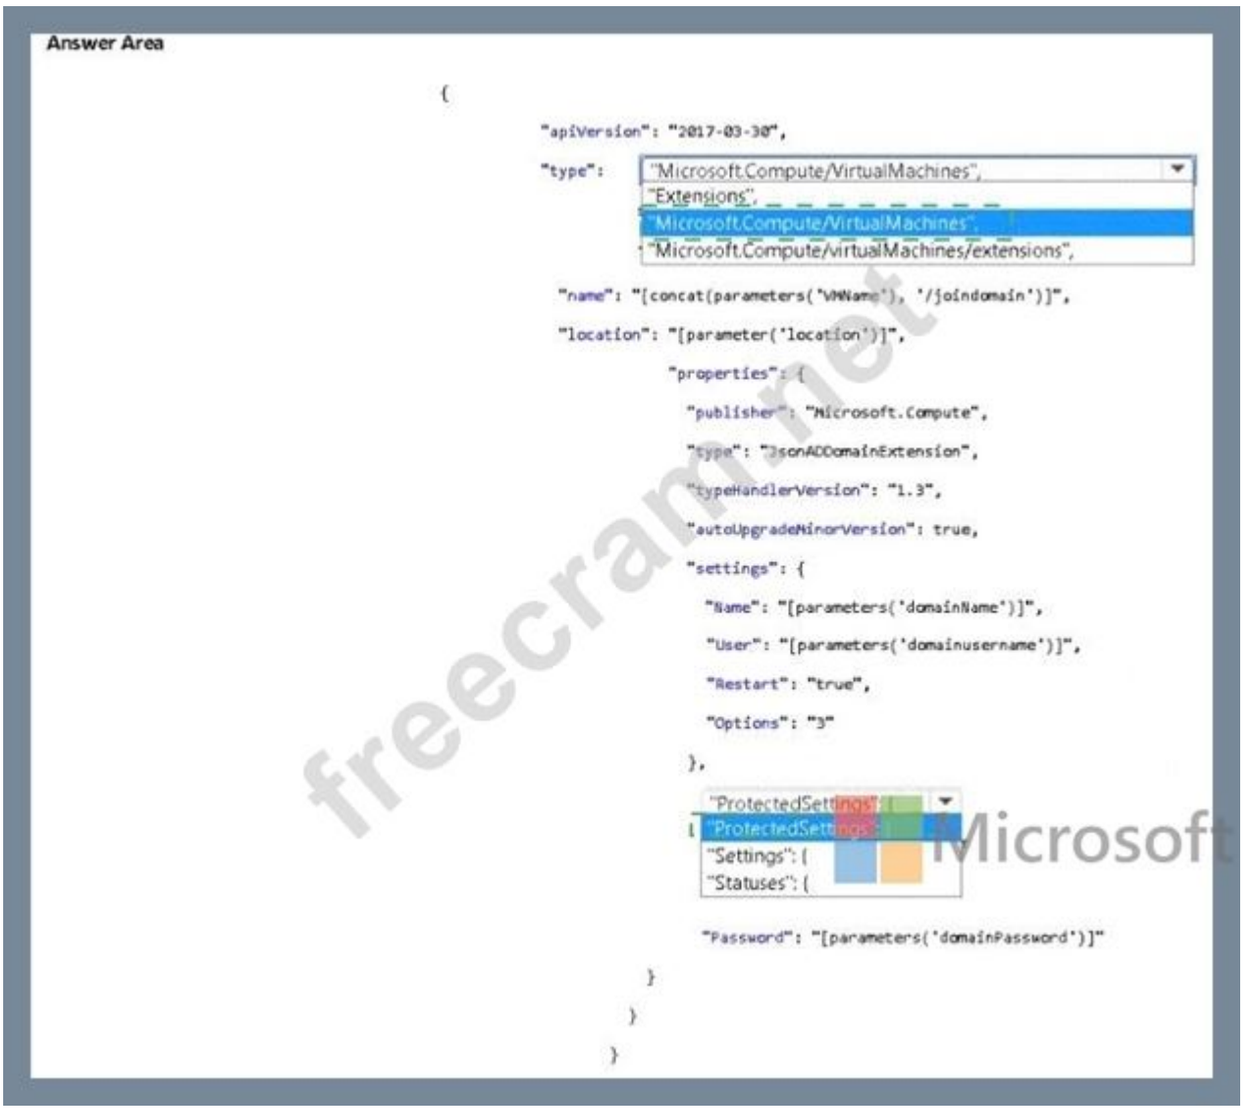

Explanation:

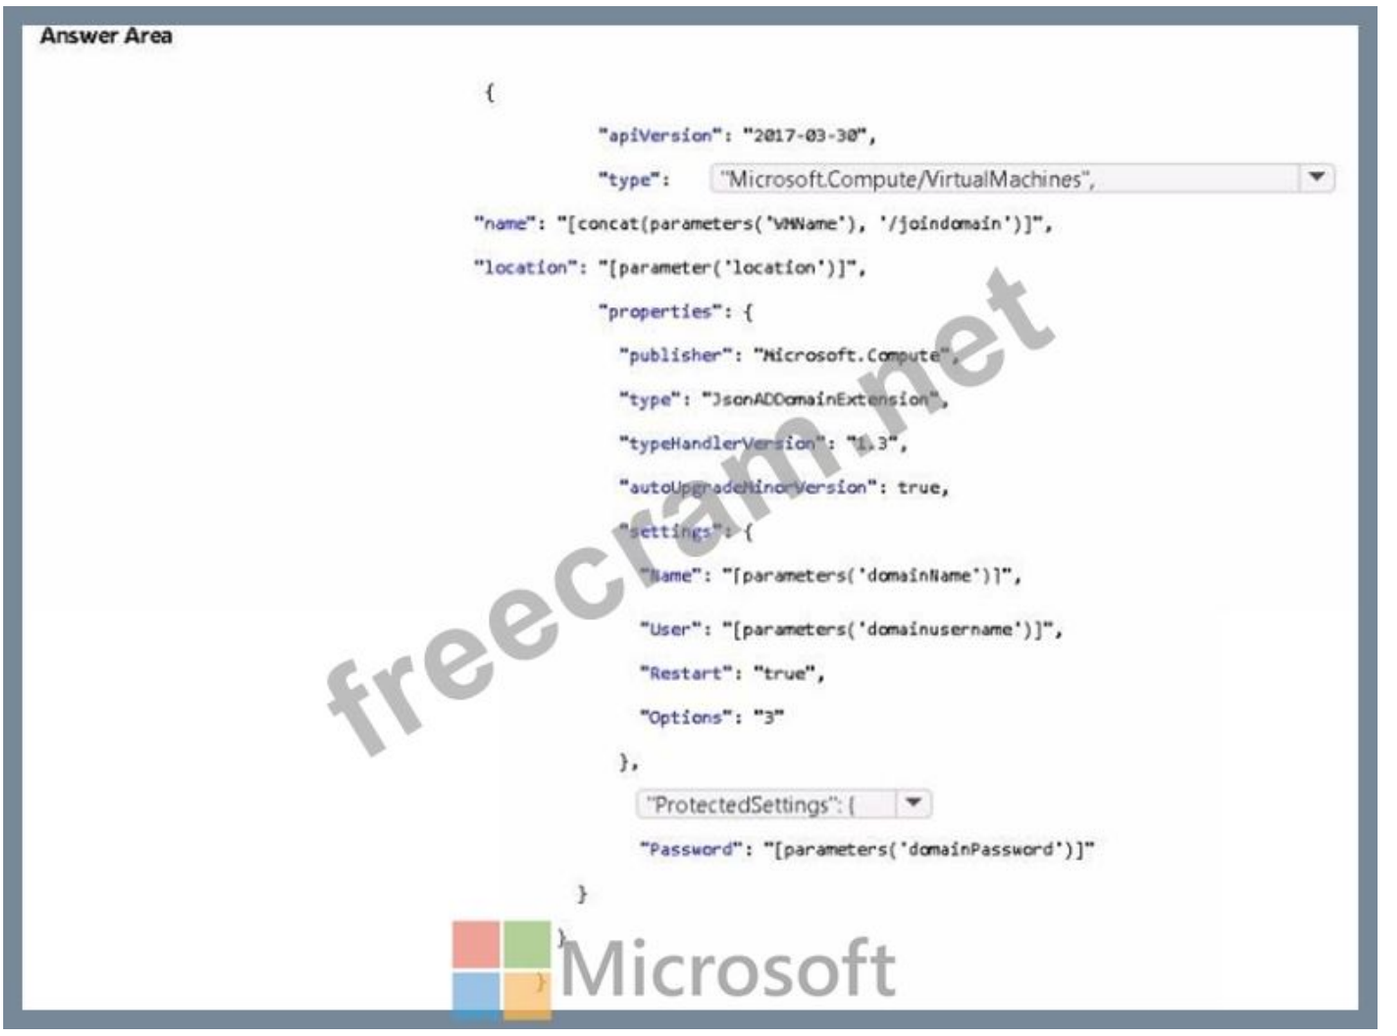

In Azure, to automate domain join operations for virtual machines through an Azure Resource Manager (ARM) template, you must configure the JsonADDomainExtension-a VM extension provided by Microsoft.

This extension allows a Windows virtual machine to be automatically joined to an Active Directory domain after deployment.

According to the Microsoft Azure documentation ("JsonADDomainExtension for Windows"), the extension must be defined under the Microsoft.Compute/virtualMachines/extensions resource type, because it is a child resource of a specific VM. This structure allows Azure to attach and configure the extension directly to the virtual machine during or after deployment.

The correct syntax within the ARM template must include:

* Type: "Microsoft.Compute/virtualMachines/extensions" # Specifies the extension resource type.
* Publisher: "Microsoft.Compute" # Specifies the vendor of the extension.
* Type: "JsonADDomainExtension" # Specifies the type of extension.
* ProtectedSettings: used to securely store sensitive information like domain passwords (encrypted automatically by Azure).

The ProtectedSettings section ensures credentials such as the domain password remain confidential, while Settings contain non-sensitive configuration parameters (domain name, username, etc.).

This configuration ensures the VM automatically joins the target domain securely upon creation.

Final Verified Answer:
* Type: "Microsoft.Compute/virtualMachines/extensions"
* ProtectedSettings: "ProtectedSettings"

NEW QUESTION: 27
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Microsoft Entra tenant named Adatum.com and an Azure Subscription named Subscription1.
Adatum.com contains a group named Developers. Subscription1 contains a resource group named Dev.
You need to provide the Developers group with the ability to create Azure Logic Apps in the Dev resource group.
Solution: On Subscription1, you assign the DevTest Labs User role to the Developers group.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
The DevTest Labs User role is a built-in Azure role designed specifically to support operations within Azure DevTest Labs. According to the Microsoft Azure Administrator documentation on built-in roles, this role permits users to connect to, start, stop, restart, and use virtual machines inside a DevTest Lab, but it does not grant permissions to create or manage Azure resources outside of DevTest Labs, such as Azure Logic Apps.
In this scenario, the requirement is to allow the Developers group to create Azure Logic Apps in the Dev resource group. Azure Logic Apps are managed through the Microsoft.Logic resource provider and require permissions such as Microsoft.Logic/workflows/write. These permissions are included in roles like Logic App Contributor or the more general Contributor role at the resource group scope.
The DevTest Labs User role does not include permissions to deploy or manage Logic Apps, even when assigned at the subscription level. Therefore, assigning this role does not meet the stated goal. Microsoft documentation explicitly states that DevTest Labs User is limited to DevTest Lab resources and cannot be used as a general deployment role.
Final Answer: B. No

NEW QUESTION: 28
You have an Azure subscription named Subscription1 that contains a virtual network named VNet1. VNet1 is in a resource group named RG1.
Subscription1 has a user named User1. User1 has the following roles;

* Reader
* Security Admin
* Security Reader

You need to ensure that User1 can assign the Reader role for VNet1 to other users. What should you do?

A. Assign User1 the Contributor role for VNet1.
B. Remove User from the Security Reader and Reader roles tot Subscription1.
C. Assign User1 the Network Contributor role for VNet1.
D. Assign User1 the User Access Administrator role for VNet1

Answer: (SHOW ANSWER)

https://docs.microsoft.com/en-us/azure/role-based-access-control/rbac-and-directory-admin-roles#:~:text=The%20User%20Access%20Administrator%20role%20enables%20the%20user%20to%20grant,Azure%20subscriptions%20and%20management%20groups.

NEW QUESTION: 29

You have an Azure virtual machine named VM1 and a Recovery Services vault named Vault1. You create a backup Policy1 as shown in the exhibit. (Click the Exhibit tab.)

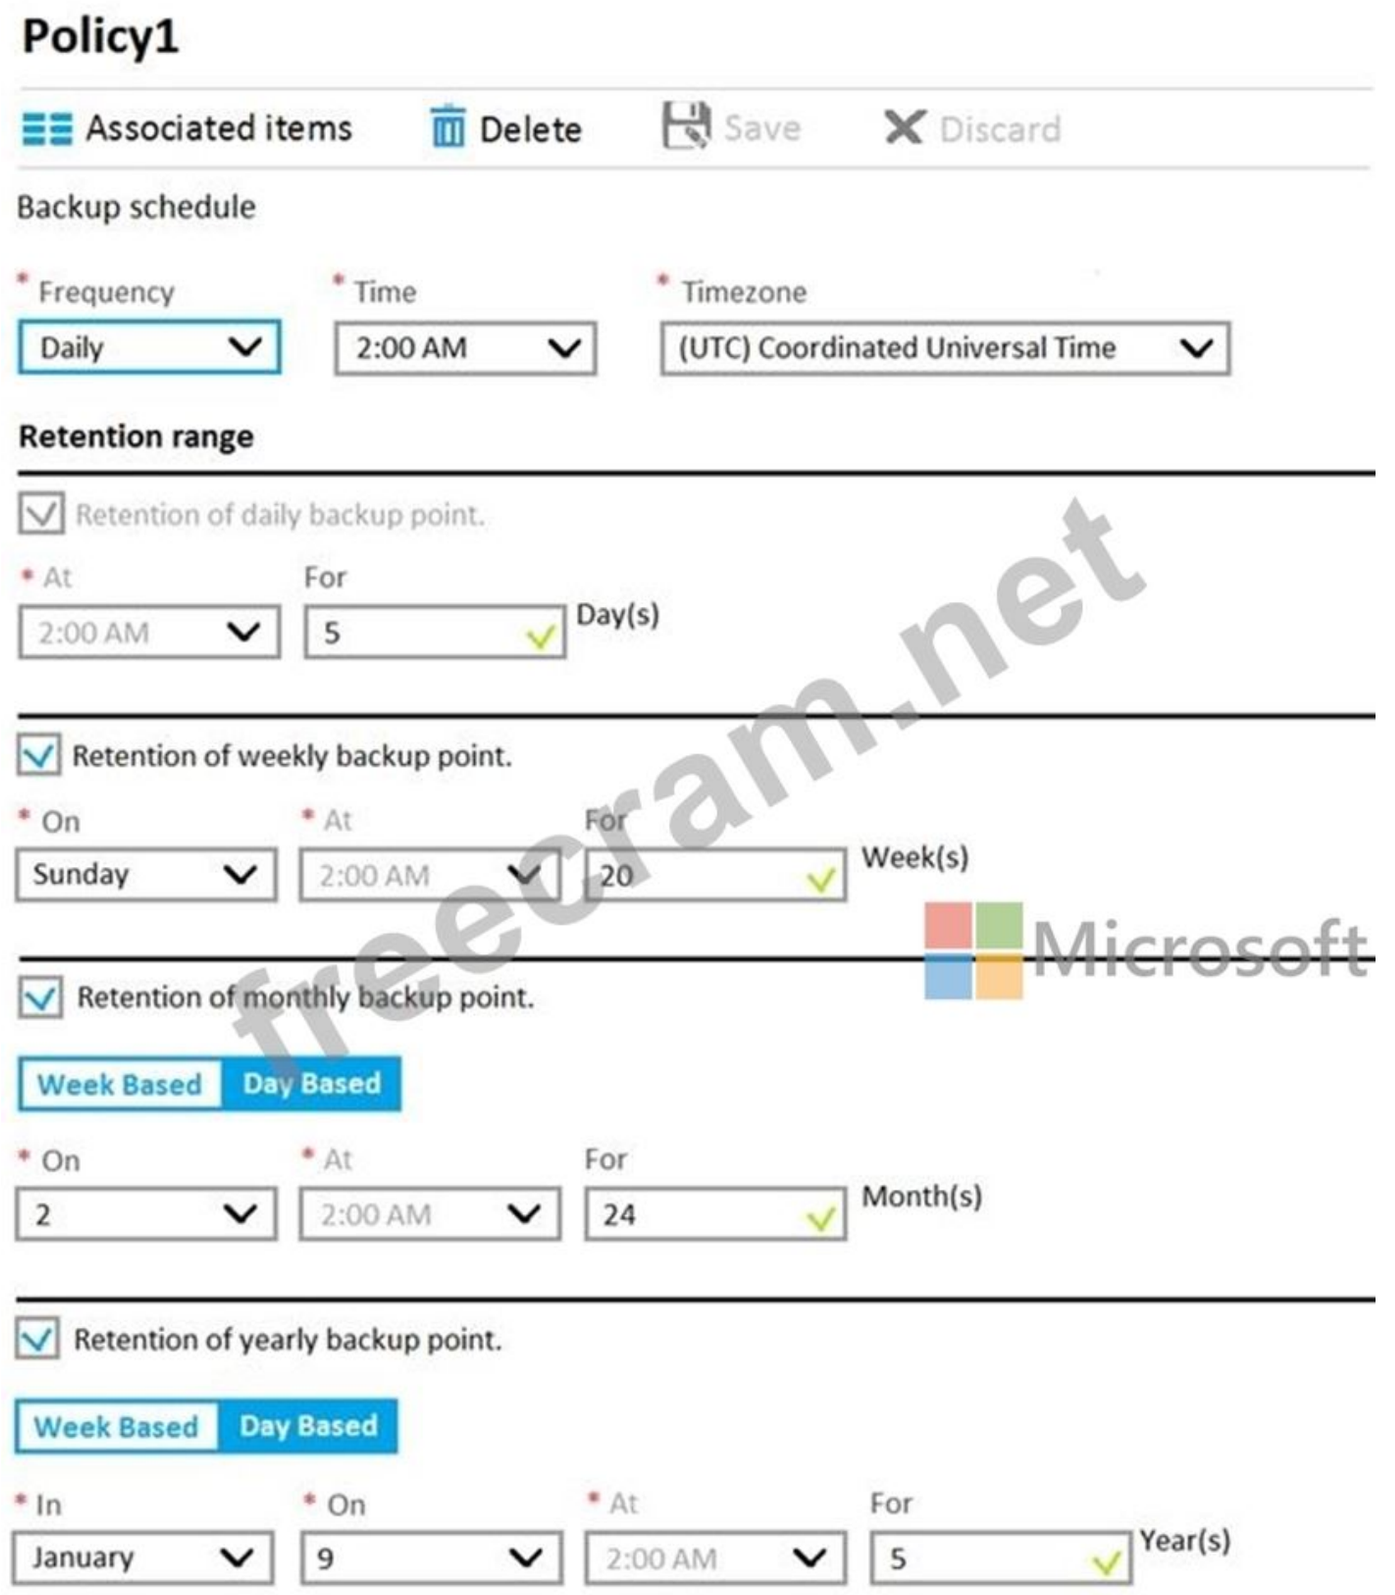

You configure the backup of VM1 to use Policy1 on Thursday, January 1.

You need to identify the number of available recovery points for VM1.

How many recovery points are available on January 8 and on January 15? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

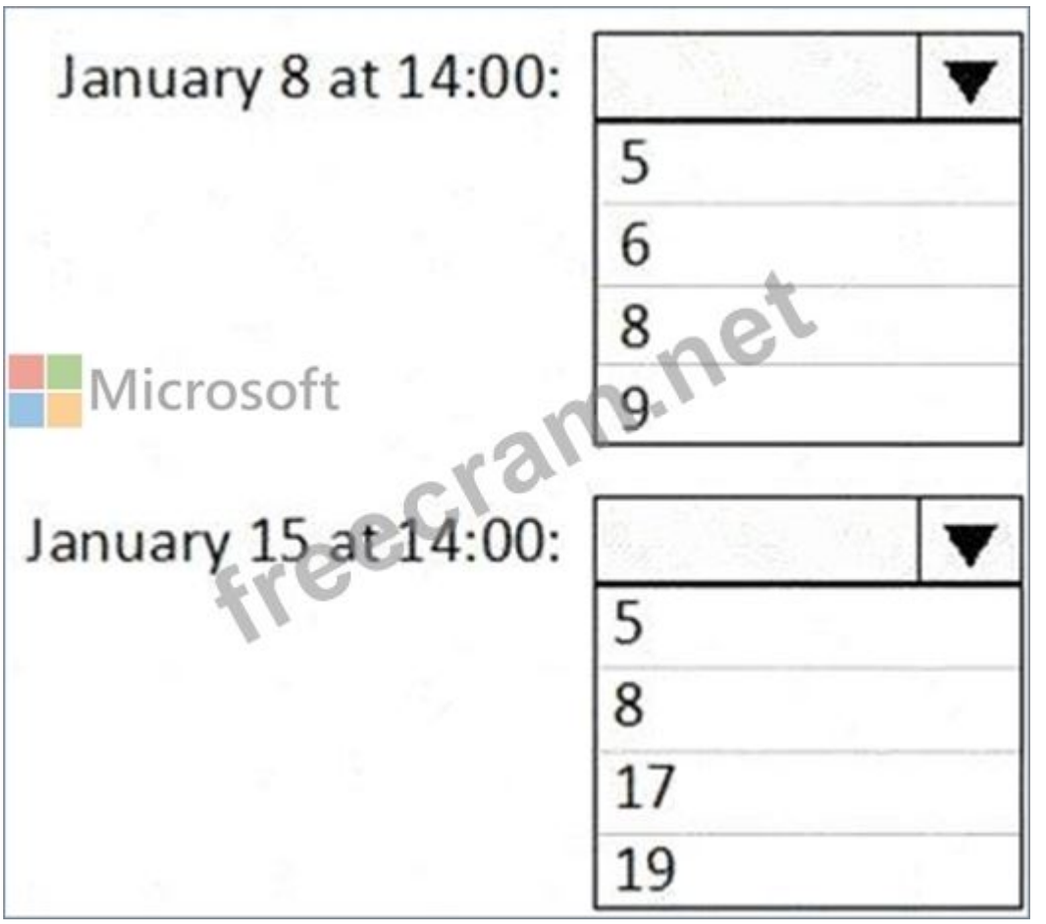

Answer:

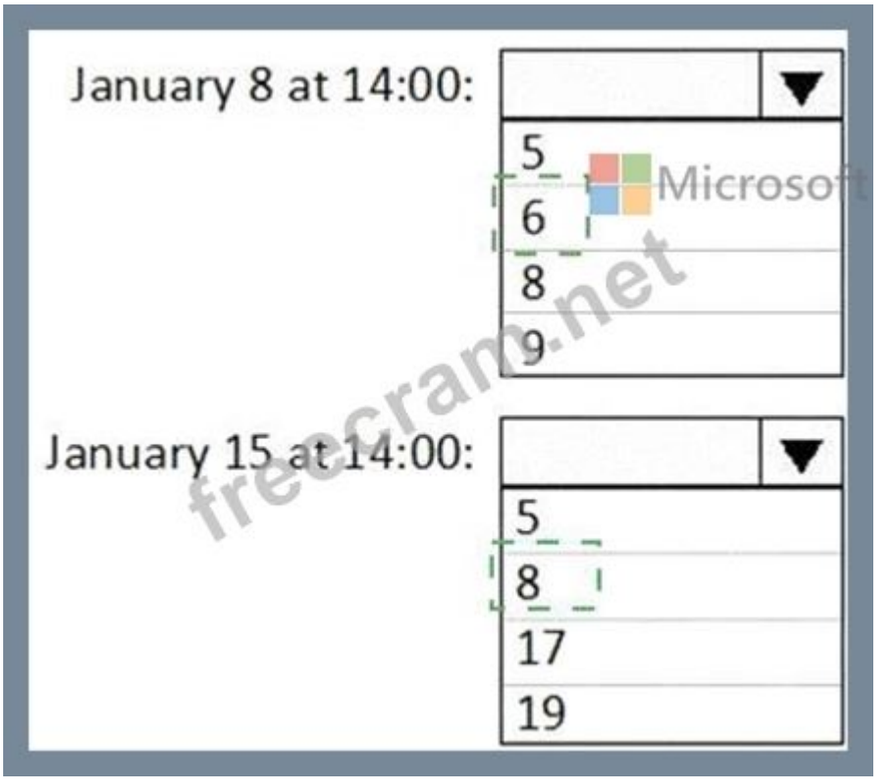

Explanation:

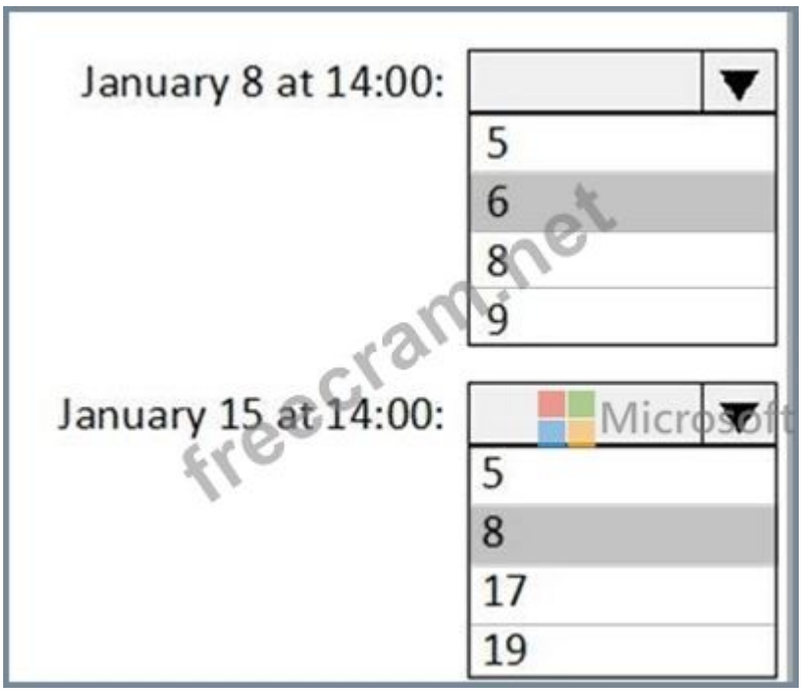

NEW QUESTION: 30

You have an Azure DNS zone named adatum.com. You need to delegate a subdomain named research.adatum.
com to a different DNS server in Azure. What should you do?

A. Create an PTR record named research in the adatum.com zone.

B. Create an NS record named research in the adatum.com zone.

C. Modify the SOA record of adatum.com.

D. Create an A record named *. research in the adatum.com zone

Answer: (SHOW ANSWER)

In Azure DNS, delegating a subdomain (such as research.adatum.com) to another DNS server or DNS zone requires creating a Name Server (NS) record in the parent domain (adatum.com).

Here's how DNS delegation works according to the Microsoft Azure DNS documentation:

* The parent DNS zone (adatum.com) must contain an NS record set for the subdomain (research).
* This NS record points to the authoritative name servers of the child DNS zone (research.adatum.com).
* When a DNS query is made for any record under research.adatum.com, the DNS resolution process will follow the delegation and forward the query to the specified DNS servers.

Example configuration:

Record Type

Name

Value

NS

research

ns1-xx.azure-dns.com

ns2-xx.azure-dns.net

This delegation ensures that the subdomain research.adatum.com is managed independently, possibly by another team or subscription.
Other options:
* PTR record: Used for reverse DNS lookups, not for delegation.
* SOA record: Defines start of authority for a zone, not delegation.
* A record: Maps a name to an IP address, not for delegating zones.
# Final Verified Answer: B. Create an NS record named research in the adatum.com zone

NEW QUESTION: 31
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You need to ensure that an Azure Active Directory (Azure AD) user named Admin1 is assigned the required role to enable Traffic Analytics for an Azure subscription.
Solution: You assign the Owner role at the subscription level to Admin1.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
To enable Traffic Analytics for an Azure subscription, the user must have sufficient privileges to configure Network Watcher, NSG flow logs, and the associated Log Analytics workspace.
As per Microsoft Azure documentation, the following built-in roles can enable Traffic Analytics:
* Owner
* Contributor
* Reader
* Network Contributor
The Owner role provides full access to all resources, including the right to delegate permissions and modify configurations. Since the Owner role includes complete management capabilities for all Azure resources at the subscription level, this role absolutely meets the requirements for enabling Traffic Analytics.
The Azure Network Watcher documentation clearly states:
"To enable Traffic Analytics, your account must have any one of the following roles at the subscription scope: Owner, Contributor, Reader, or Network Contributor." Therefore, assigning the Owner role to Admin1 at the subscription level ensures Admin1 has the required permissions to enable Traffic Analytics.

NEW QUESTION: 32
You need to define a custom domain name for Azure AD to support the planned infrastructure.
Which domain name should you use?
A. ad.humongousinsurance.com
B. humongousinsurance.onmicrosoft.com
C. humongousinsurance.local
D. humongousinsurance.com
Answer: (SHOW ANSWER)
Every Azure AD directory comes with an initial domain name in the form of domainname.onmicrosoft.com.
The initial domain name cannot be changed or deleted, but you can add your corporate domain name to Azure AD as well. For example, your organization probably has other domain names used to do business and users who sign in using your corporate domain name. Adding custom domain names to Azure AD allows you to assign user names in the directory that are familiar to your users, such as 'alice@contoso.com.' instead of 'alice@domain name.onmicrosoft.com'.
Scenario:
Network Infrastructure: Each office has a local data center that contains all the servers for that office. Each office has a dedicated connection to the Internet.
Humongous Insurance has a single-domain Active Directory forest named humongousinsurance.com Planned Azure AD Infrastructure: The on-premises Active Directory domain will be synchronized to Azure AD.
References:
https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/add-custom-domain

NEW QUESTION: 33
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You manage a virtual network named VNet1 that is hosted in the West US Azure region.
VNet1 hosts two virtual machines named VM1 and VM2 that run Windows Server.

You need to inspect all the network traffic from VM1 to VM2 for a period of three hours.
Solution: From Azure Monitor, you create a metric on Network in and Network Out.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
Azure Monitor metrics such as Network In and Network Out only provide aggregated performance data - i.
e., the total volume of bytes entering or leaving a network interface card (NIC). These metrics are useful for monitoring throughput, trends, or network utilization but do not capture or inspect packet-level data between two virtual machines.
To inspect network traffic between two Azure VMs (such as between VM1 and VM2), you must use Azure Network Watcher's Packet Capture or Connection Monitor feature.
According to the Microsoft Azure Administrator study guide and official documentation:
"Use Azure Network Watcher to monitor, diagnose, and view metrics for resources in a virtual network. The Packet Capture feature enables you to capture network traffic to and from a virtual machine and save it to an Azure Storage account for analysis." Therefore, simply creating metrics in Azure Monitor on Network In/Out will not meet the goal, because metrics do not provide detailed traffic inspection, only summary statistics.
To fulfill the requirement ("inspect all the network traffic from VM1 to VM2 for three hours"), the correct approach would be to use Network Watcher # Packet Capture, specifying VM1 as the target, capturing outbound traffic to VM2's IP, and setting the session duration for 3 hours.
Hence, the given solution does not meet the goal.

NEW QUESTION: 34
You have an Azure subscription that contains an Azure virtual machine named VM1. VM1 runs a financial reporting app named App1 that does not support multiple active instances.
At the end of each month, CPU usage for VM1 peaks when App1 runs.
You need to create a scheduled runbook to increase the processor performance of VM1 at the end of each month.
What task should you include in the runbook?
A. Add the Azure Performance Diagnostics agent to VM1.
B. Modify the VM size property of VM1.
C. Add VM1 to a scale set.
D. Increase the vCPU quota for the subscription.
E. Add a Desired State Configuration (DSC) extension to VM1.
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: 35

You have an Azure subscription that contains a user named User1 and the resources shown in the following table.

Name	Type
RG1	Resource group
networkinterface1	Virtual network interface
NSG1	Network security group (NSG)

NSG1 is associated to networkinterface1.

User1 has role assignments tor NSG1 as shown in the following table.

Role	Scope
Contributor	This resource
Reader	Subscription (Inherited)
Storage Account Contributor	Resource group (Inherited)

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE Each correct selection is worth one point.

Answer Area

Statements	Yes	No
User1 can create a storage account in RG1.	○	○
User1 can modify the DNS settings of networkinterface1.	○	○
User1 can create an inbound security rule to filter inbound traffic to networkinterface1.	○	○

Answer:

Answer Area

Statements	Yes	No
User1 can create a storage account in RG1.	○	○
User1 can modify the DNS settings of networkinterface1.	○	○
User1 can create an inbound security rule to filter inbound traffic to networkinterface1.	○	○

Explanation:

Answer Area

Statements	Yes	No
User1 can create a storage account in RG1.	◉	○
User1 can modify the DNS settings of networkinterface1.	○	◉
User1 can create an inbound security rule to filter inbound traffic to networkinterface1.	◉	○

NEW QUESTION: 36

You have an Azure subscription that contains a virtual network named VNET in the East Us 2 region. A network interface named VM1-NI is connected to VNET1.

You successfully deploy the following Azure Resource Manager template.

```
{
    "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')]"
      ],
      "storageProfile": {
        "imageReference": "[variables('image')]",
        "osDisk": {
          "createOption": "FromImage"
        }
      },
      "networkProfile": {
        "networkInterfaces": [
          {
```

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: 37

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure virtual machine named VM1. VM1 was deployed by using a custom Azure Resource Manager template named ARM1.json.
You receive a notification that VM1 will be affected by maintenance.
You need to move VM1 to a different host immediately.
Solution: From the Overview blade, you move the virtual machine to a different subscription.
Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

When Microsoft schedules maintenance for an Azure virtual machine (VM), the affected VM may experience a short downtime or host migration. In some cases, administrators can proactively move the VM to a new physical host before maintenance begins to minimize disruption.

According to the Microsoft Azure Administrator documentation, the only supported method to proactively move a VM to a new host is to use the Redeploy feature available under VM # Redeploy + reapply in the Azure portal or through PowerShell (Set-AzVM -Redeploy).

The Redeploy operation stops the VM, moves it to a new physical host within the same region, and then restarts it automatically. This ensures the VM is placed on healthy hardware immediately, mitigating any pending maintenance impact.

In contrast, moving a VM to a different subscription:

* Is a metadata-level operation, not a host-level migration.
* Requires the VM to be in a deallocated state before the move.
* Does not cause the VM to change its underlying physical host.
* Is typically used for billing or management reorganization, not to mitigate maintenance events.

Therefore, moving the VM to another subscription does not achieve the goal of moving VM1 to a different physical host. The VM remains on the same compute node until redeployed or maintenance is automatically performed by Azure.

The correct solution in this scenario is to use "Redeploy" from the VM blade, which explicitly triggers host relocation.

NEW QUESTION: 38

You have an Azure Active Directory (Azure AD) tenant that has the initial domain name.

You have a domain name of contoso.com registered at a third-party registrar.

You need to ensure that you can create Azure AD users that have names containing a suffix of @contoso.com.

Which three actions should you perform in sequence? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

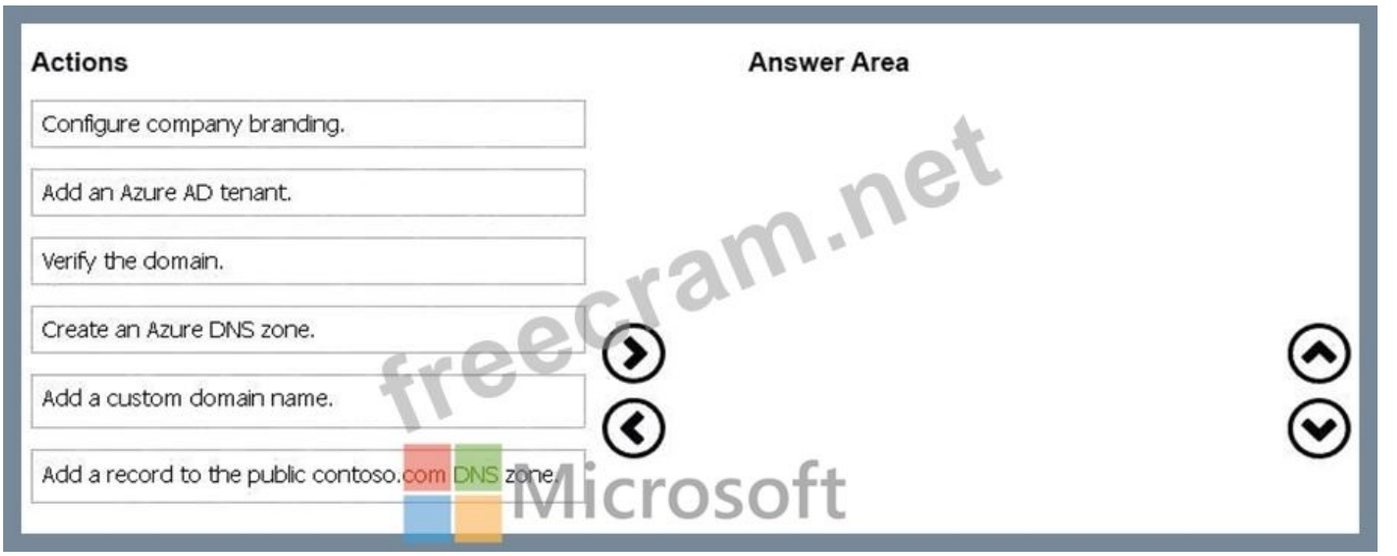

Answer:

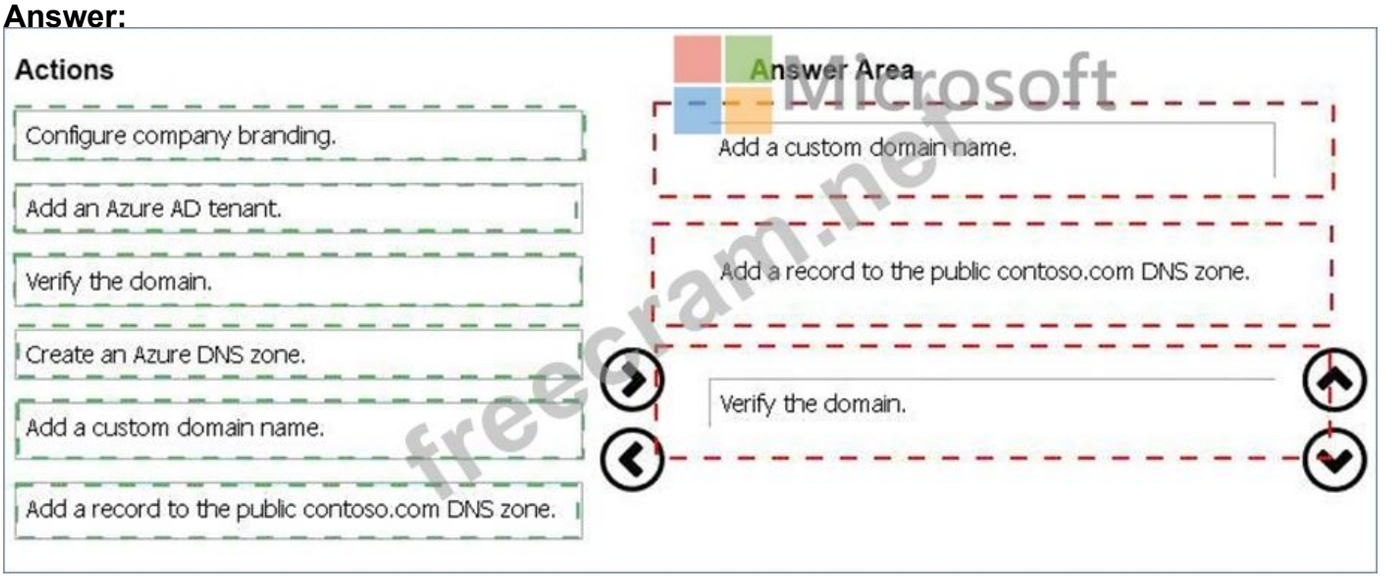

Explanation:

The process is simple:

Add the custom domain name to your directory

Add a DNS entry for the domain name at the domain name registrar

Verify the custom domain name in Azure AD

References: https://docs.microsoft.com/en-us/azure/dns/dns-web-sites-custom-domain

NEW QUESTION: 39

You have the Azure resources shown on the following exhibit.

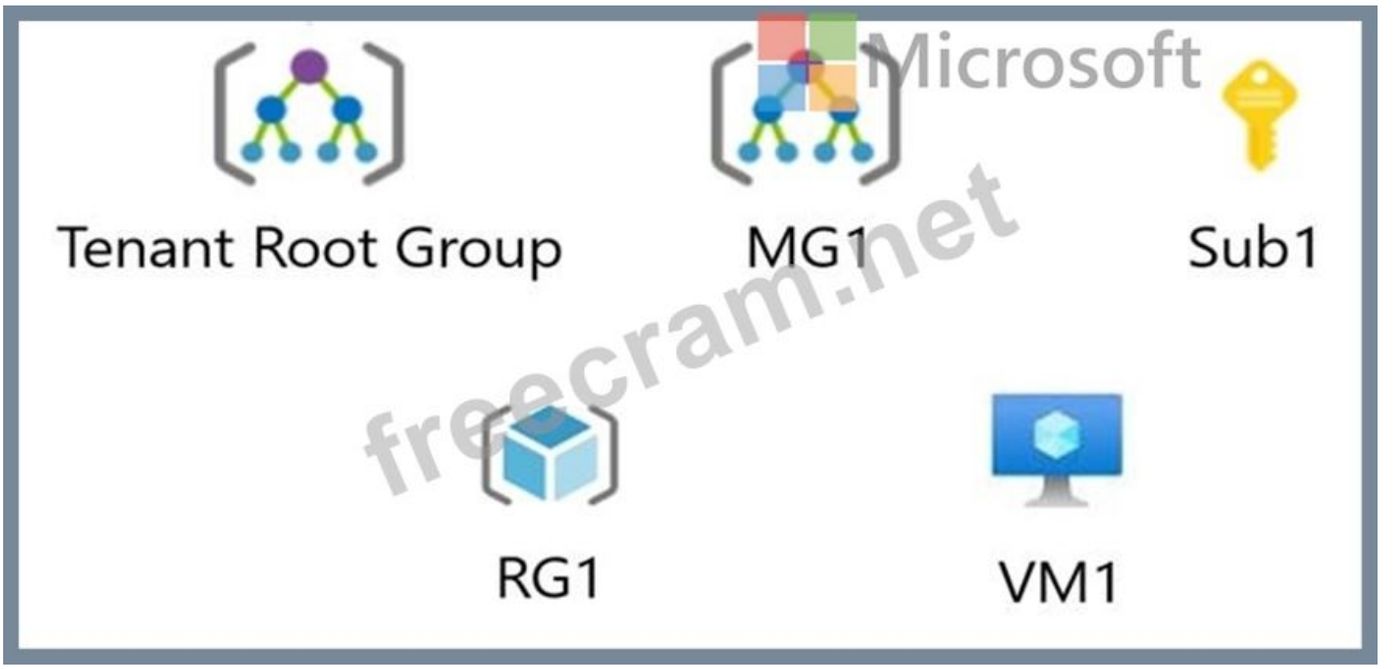

You plan to track resource usage and prevent the deletion of resources.

To which resources can you apply locks and tags? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Locks:

- RG1 and VM1 only
- Sub1 and RG1 only
- Sub1, RG1, and VM1 only
- MG1, Sub1, RG1, and VM1 only
- Tenant Root Group, MG1, Sub1, RG1, and VM1

Tags:

- RG1 and VM1 only
- Sub1 and RG1 only
- Sub1, RG1, and VM1 only
- MG1, Sub1, RG1, and VM1 only
- Tenant Root Group, MG1, Sub1, RG1, and VM1

Answer:

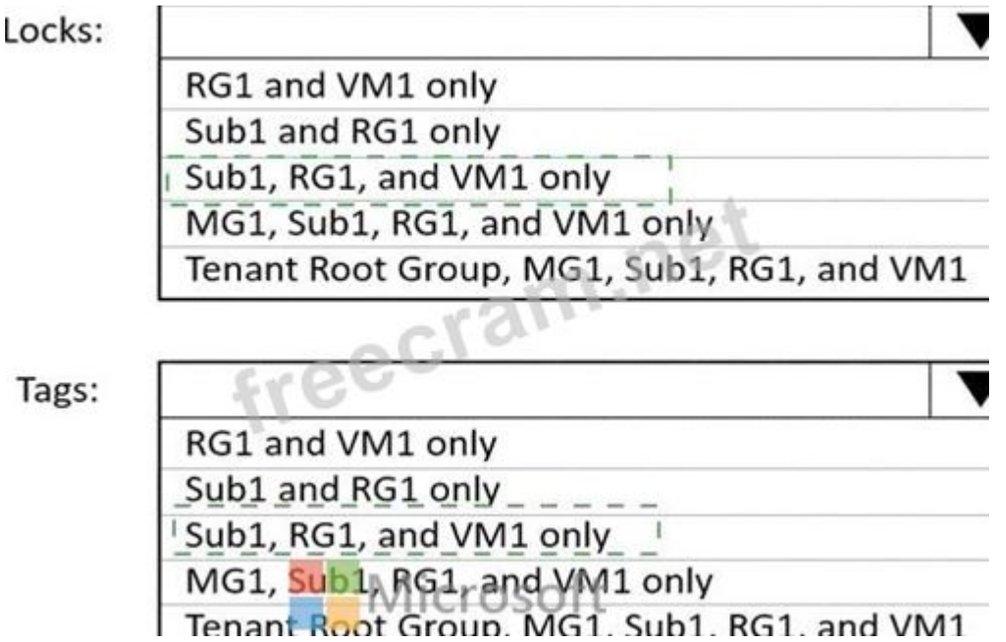

Explanation:

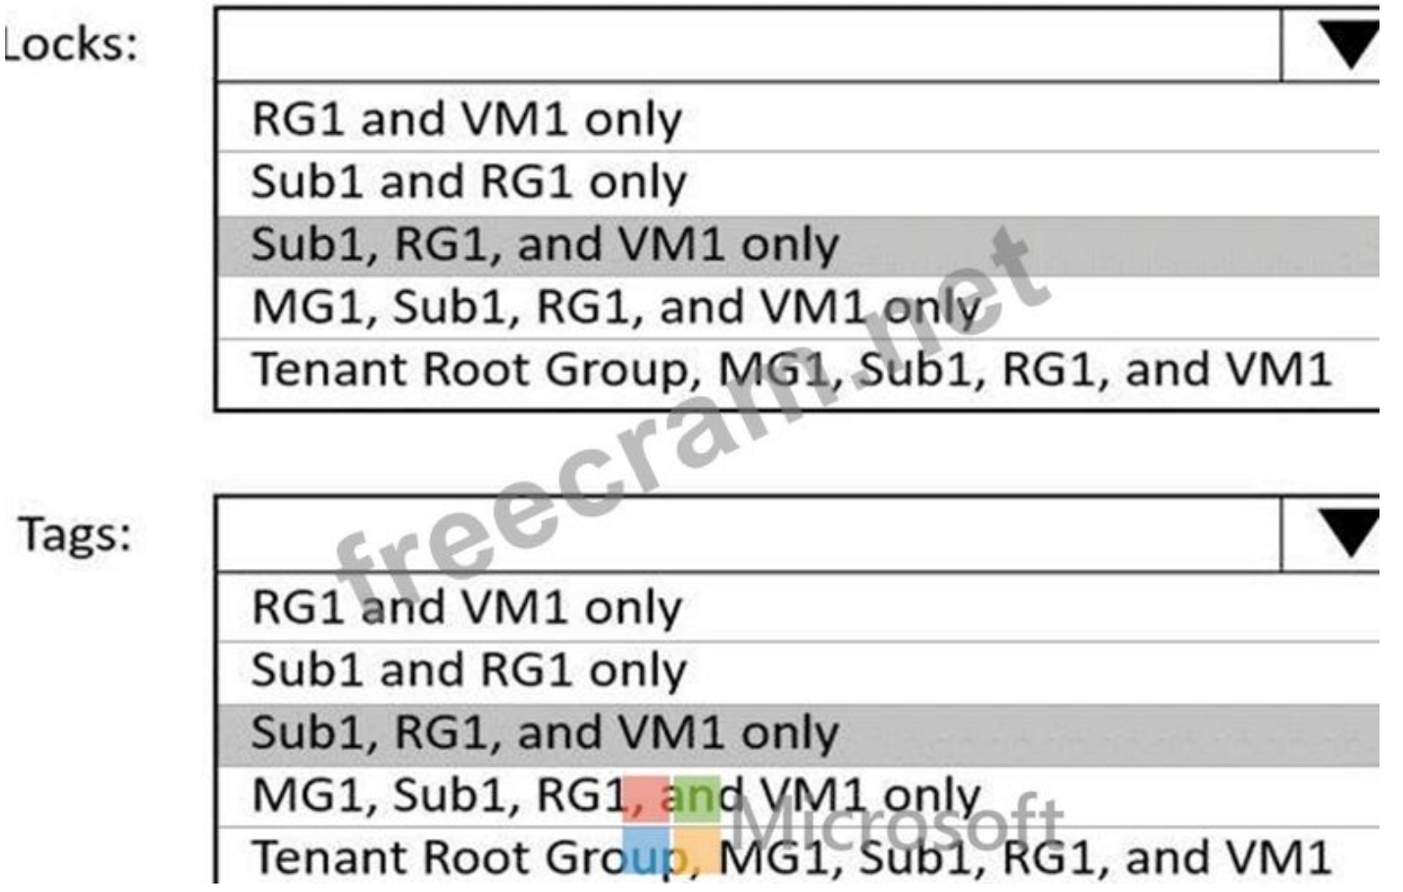

Box 1: Sub1, RG1, and VM1 only

You can lock a subscription, resource group, or resource to prevent other users in your organization from accidentally deleting or modifying critical resources.

Box 2: Sub1, RG1, and VM1 only

You apply tags to your Azure resources, resource groups, and subscriptions.

Reference:

https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/lock-resources?

tabs=json
https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/tag-resources?tabs=json

NEW QUESTION: 40

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure subscription that contains 10 virtual networks. The virtual networks are hosted in separate resource groups.
Another administrator plans to create several network security groups (NSGs) in the subscription.
You need to ensure that when an NSG is created, it automatically blocks TCP port 8080 between the virtual networks.
Solution: You configure a custom policy definition, and then you assign the policy to the subscription.
Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

A custom policy definition is a way to define your own rules for using Azure resources. You can use custom policies to enforce compliance, security, cost management, or organization-specific requirements. However, a custom policy definition alone is not enough to meet the goal of automatically blocking TCP port 8080 between the virtual networks. You also need to create a policy assignment that applies the custom policy definition to the scope of the subscription. A policy assignment is the link between a policy definition and an Azure resource. Without a policy assignment, the custom policy definition will not take effect. Therefore, the solution does not meet the goal.
References:
Tutorial: Create a custom policy definition
Create and manage policies to enforce compliance

NEW QUESTION: 41

You have an Azure subscription that contains virtual machine named VM1.
You need to back up VM. The solution must ensure that backups are stored across three availability zones in the primary region.
Which three actions sh ould you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

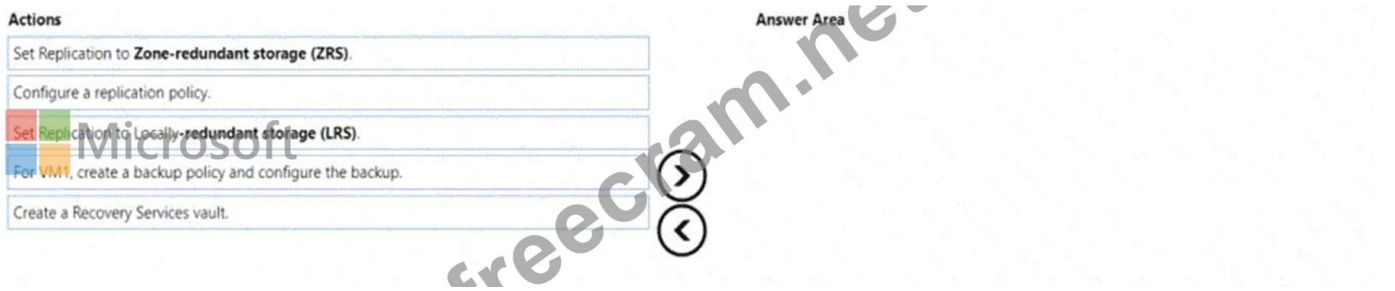

Answer:

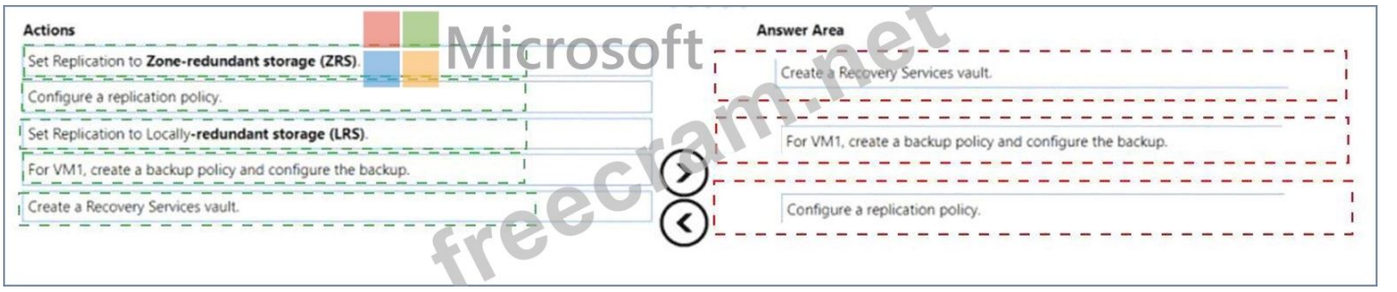

Explanation:

According to 1, Availability Zones are unique physical locations within an Azure region that provide high availability and disaster recovery for your virtual machines. To back up your VM across three availability zones in the primary region, you need to perform the following actions in sequence:

Create a Recovery Services vault2 that will store your backups and enable geo-redundancy for cross-region protection.

For VM1, create a backup policy and configure the backup2 to use the Recovery Services vault as the backup destination.

Configure a replication policy1 that will replicate your VM1 to another availability zone in the same region.

NEW QUESTION: 42

You have an Azure subscription that hat Traffic Analytics configured.

You deploy a new virtual machine named VM1 that has the following settings:

* Region- East US
* Virtual network: VNet1
* NIC network security group: NSG1

You need to monitor VM1 traffic by using Traffic Analytics.

Which settings should you configure?

A. Diagnostic settings for VM1

B. Insights for VM1

C. NSG flow logs for NSG1

D. Diagnostic settings for NSG1

Answer: (SHOW ANSWER)

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: 43

You have two subscriptions named Subscription1 and Subscription2. Each subscription is associated to a different Azure AD tenant.

Subscription1 contains a virtual network named VNet1. VNet1 contains an Azure virtual machine named VM1 and has an IP address space of 10.0.0.0/16.

Subscription2 contains a virtual network named VNet2. VNet2 contains an Azure virtual machine named VM2 and has an IP address space of 10.10.0.0/24.

You need to connect VNet1 to VNet2.

What should you do first?

A. Provision virtual network gateways.

B. Modify the IP address space of VNet2.

C. Move VM1 to Subscription2.

D. Move VNet1 to Subscription2.

Answer: A (LEAVE A REPLY)

Azure VNet peering is the most common way to connect virtual networks, but in this scenario, VNet1 and VNet2 exist in different Azure AD tenants. Global or cross-tenant peering requires that both VNets be within the same Azure Active Directory tenant.

Because the question specifies that each subscription is associated with a different Azure AD tenant, VNet peering is not supported across tenants.

To establish connectivity between VNets in different tenants, the correct solution is to configure VPN gateways in both virtual networks and create a VNet-to-VNet VPN connection.

This type of connection uses the Azure VPN Gateway service to securely tunnel traffic between the VNets using IPsec/IKE protocols over the Internet.

Changing IP address spaces or moving VNets/subscriptions does not solve the cross-tenant issue securely or efficiently.

Hence, before connecting VNet1 and VNet2, you must provision virtual network gateways in both VNets.

NEW QUESTION: 44

Which blade should you instruct the finance department auditors to use?

A. invoices

B. partner information

C. cost analysis

D. External services

Answer: (SHOW ANSWER)

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.

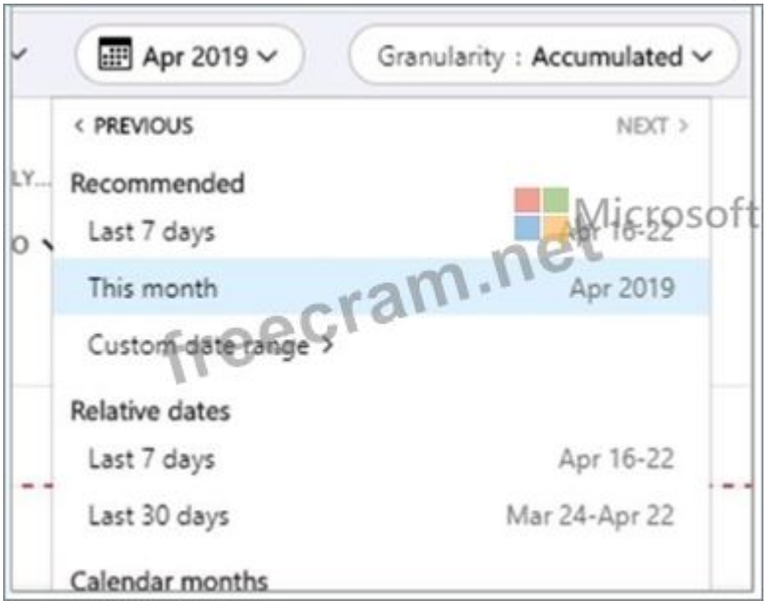

Invoice: Incorrect Option

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: 45

You need to configure encryption for the virtual machines. The solution must meet the technical requirements.

Which virtual machines can you encrypt?

A. VM1 and VM3

B. VM2 and VM3

C. VM2 and VM4

D. VM4 and VM5

Answer: (SHOW ANSWER)

To determine which virtual machines can be encrypted, we must refer to the technical requirement:

"Whenever possible, use Azure Disk Encryption and a key encryption key (KEK) to encrypt the virtual machines." According to the Microsoft Azure Administrator documentation, Azure Disk Encryption (ADE) uses BitLocker for Windows virtual machines and DM-Crypt for Linux virtual

machines to encrypt OS and data disks. However, not all VM types and disk configurations are supported for ADE.
From the provided configuration:
VM
OS
Description
Disk Type
VM1
RHEL
Uses ephemeral OS disks
Not supported
VM2
Windows Server 2022
Has a basic volume
Supported
VM3
RHEL
Uses standard SSDs
Supported by DM-Crypt, but not optimal
VM4
Windows Server 2022
Uses Write Accelerator disks
Supported
VM5
Windows Server 2022
Has a dynamic volume
Not supported by ADE
Step-by-step Analysis (Based on Microsoft Docs):
Ephemeral OS disks (VM1):
These are not compatible with Azure Disk Encryption because they are stored on local temporary storage and are not persisted to Azure Storage.
"Ephemeral OS disks cannot be encrypted using Azure Disk Encryption because they reside on local VM storage." - [Microsoft Learn: Azure Disk Encryption prerequisites] Dynamic volumes (VM5):
Azure Disk Encryption does not support dynamic disks - only basic disks are supported.
"Azure Disk Encryption does not support dynamic disks; only basic disks can be encrypted." - [Microsoft Learn: ADE limitations] Write Accelerator disks (VM4):
These disks can be encrypted if they are standard OS/data disks with Write Accelerator enabled.
Microsoft confirms that ADE supports Write Accelerator-enabled disks on supported VM sizes.
Linux VMs (VM3) with standard SSDs can use DM-Crypt encryption, but in this question, the requirement specifies to use Azure Disk Encryption with a Key Encryption Key (KEK) - KEKs are

supported only for Windows and Linux VMs that use managed disks and not ephemeral disks. However, VM3 uses standard SSDs (supported by ADE) and VM2 and VM4 meet the technical requirement of using ADE with KEK, but Azure prefers Windows VMs for KEK integration because of BitLocker integration and Key Vault support.
Therefore, the verified supported VMs for Azure Disk Encryption with KEK are:
VM2 (Windows Server 2022, basic volume)
VM4 (Windows Server 2022, Write Accelerator disks)
Supporting Microsoft Documentation (Extracted):
"Azure Disk Encryption helps protect and safeguard your data to meet your organizational security and compliance commitments. It uses BitLocker for Windows VMs and DM-Crypt for Linux VMs."
"ADE is not supported on ephemeral OS disks or dynamic disks."
"You can use Key Encryption Keys (KEK) from Azure Key Vault for an added layer of security." (Source: Microsoft Learn - Azure Disk Encryption Overview, Prerequisites, and Supported Configurations for Windows and Linux VMs)
Topic 1, A. Datum Corporation
Overview
A). Datum Corporation is a consulting firm that has a main office in Montreal and branch offices in Seattle and New York.
Azure Environment
A). Datum has an Azure subscription that contains three resource groups named RG1. RG2, and RG3. The subscription contains the storage accounts shown in the following table.

Name	Kind	Location	Hierarchical namespace	Container	File share
storage1	StorageV2	West US	Yes	cont1	share1
storage2	StorageV2	West US	No	cont2	share2

The subscription .contains the virtual machines shown in the following table.

Name	Size	Operating system	Description
VM1	A	Red Hat Enterprise Linux (RHEL)	Uses ephemeral OS disks
VM2	D	Windows Server 2022	Has a basic volume
VM3	B	Red Hat Enterprise Linux (RHEL)	Uses a standard SSDs
VM4	M	Windows Server 2022	Uses Write Accelerator disks
VM5	E	Windows Server 2022	Has a dynamic volume

The subscription has an Azure container registry that contains the images shown in the following table.

Name	Operating system
Image1	Windows Server
Image2	Linux

The subscription contains the resources shown in the following table.

Name	Description	In resource group
Workspace	Log Analytics workspace	RG1
WebApp1	Azure App Service web app	RG1
VNet1	Virtual network	RG2
zone1.com	Azure Private DNS zone	RG3

The subscription contains an Azure key vault named Vaultl.

Vault! contains the certificates shown in the following table.

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

Name	Type	Description
Key1	RSA	Has a key size of 4096
Key2	EC	Has Elliptic curve name set to P-256

Vaultl contains the keys shown in the following table.

Microsoft Entra Environment

A). Datum has a Microsoft Entra tenant named adatum.com that is linked to the Azure subscription and contains the users shown in the following table.

The lenant contains the groups shown in the following table.

Name	Type
Group1	Security group
Group2	Microsoft 365 group

The adatum.com tenant has a custom security attribute named Attribute1.

Planned Changes

A). Datum plans to implement the following changes:

* Configure a data collection rule {DCR) named DCR1 to collect only system events that have an event ID of

4648 from VM2 and VM4.

* In storage1, create a new container named cont2 that has the following access policies:

o Three stored access policies named Stored 1, Stored2, and Stored3 o A legal hold for immutable blob storage

* Whenever possible, use directories to organize storage account content.

* Grant User1 the permissions required to link Zone1 to VNet1.

* Assign Attribute1 to supported adatum.com resources.

* In storage2, create an encryption scope named Scope"1.

* Deploy new containers by using Image1 or Image2.

Technical Requirements

A). Datum must meet the following technical requirements:

* UseTLSforWebApp1.

* Follow the principle of least privilege.

* Grant permissions at the required scope only.
* Ensure that Scope1 is used to encrypt storage services.
* Use Azure Backup to back up cont1 and share1 as frequently as possible.
* Whenever possible, use Azure Disk Encryption and a key encryption key (KEK) to encrypt the virtual machines.

NEW QUESTION: 46
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure container registry named Registry1 that contains an image named image1.
You receive an error message when you attempt to deploy a container instance by using image1.
You need to be able to deploy a container instance by using image1.
Solution: You create a private endpoint connection for Registry1.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
This question relates to deploying an Azure Container Instance (ACI) using a container image stored in an Azure Container Registry (ACR) named Registry1. The error that occurs when deploying typically indicates authentication or network access issues between ACI and ACR.
Let's break down the scenario and solution step by step using Microsoft Learn and AZ-104 Administrator Study Guide principles.
Key Concept: How ACI Pulls Images from ACR
When you deploy a container instance that uses an image stored in an Azure Container Registry, ACI must authenticate with the registry to pull the image.
By default, ACR is private - meaning unauthenticated users (or services) cannot pull images.
Azure provides several ways to allow this communication:
Grant access using an Azure managed identity (recommended).
Enable the Admin user and provide credentials in the deployment configuration.
Assign a service principal with the AcrPull role.
Enable public network access for the registry (less secure).
What a Private Endpoint Does
Creating a private endpoint for the registry allows access to the ACR over a private IP address from within a virtual network.
According to Microsoft Learn:
"Private endpoints for Azure Container Registry enable secure access to the registry over a private link, but they do not automatically configure permissions or authentication for pulling images." While the private endpoint provides network connectivity, it does not handle

authentication or permissions.
Therefore, even with a private endpoint, ACI would still fail to pull the image unless it has appropriate role assignments or credentials to access ACR.
Why the Solution Does Not Meet the Goal
The error mentioned in the question (deployment failure) typically happens because ACI cannot authenticate to ACR.
Creating a private endpoint only provides private network connectivity, not authentication.
The correct solution would involve one of the following:
Assigning the AcrPull role to the ACI's managed identity for the ACR scope.
Enabling Admin user access on ACR and providing credentials.
Hence, simply creating a private endpoint does not resolve the deployment failure.

NEW QUESTION: 47
You have a Microsoft Entra tenant that contains the users shown in the following table.

Name	Role
Admin1	Global Administrator
Admin2	Authentication Policy Administrator
Admin3	Authentication Administrator Security Administrator

The tenant contains the groups shown in the following table.

Name	Type	Membership type
Group1	Security	Assigned
Group2	Microsoft 365	Dynamic
Group3	Mail-enabled security	Assigned

Self-service password reset (SSPR) needs to be configured for the tenant.
Which users can configure SSPR, and for which group can SSPR be enabled? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

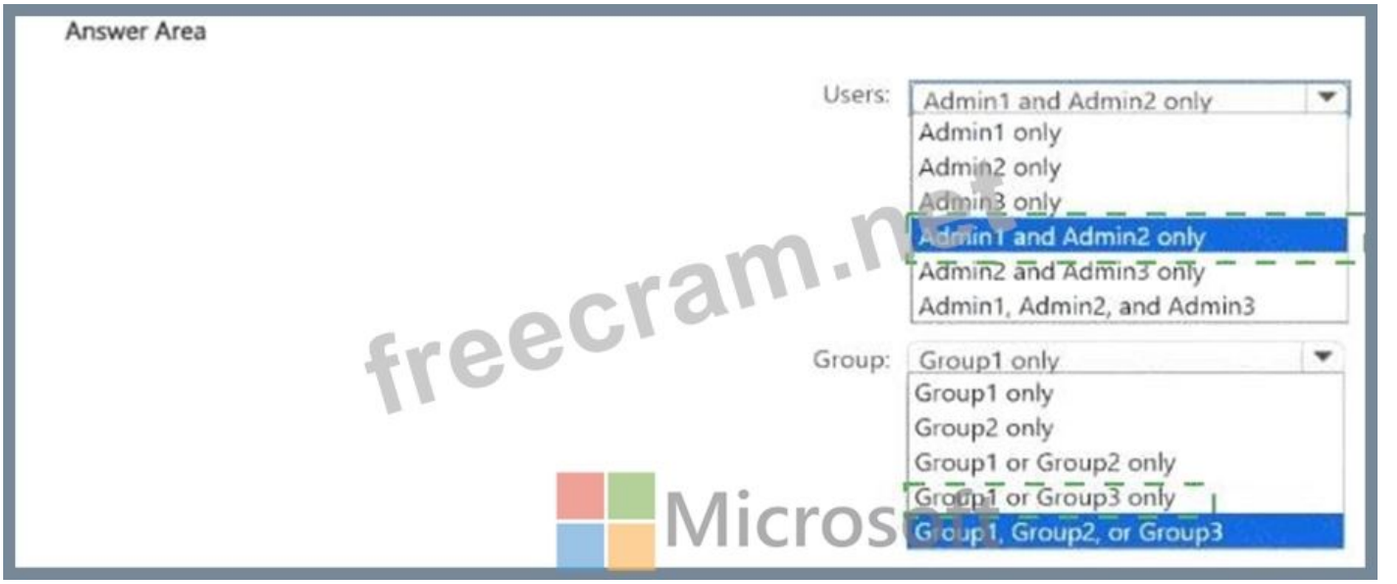

Explanation:

Users: Admin1 and Admin2 only

Group: Group1, or Group3

In Microsoft Entra ID (Azure Active Directory), Self-Service Password Reset (SSPR) is configured at the tenant level and can be scoped either to all users or to selected groups. Both who can configure SSPR and which groups are eligible are strictly defined by Microsoft Entra role permissions and group support rules.

Who can configure SSPR

Microsoft Entra documentation clearly states that only the following administrative roles are authorized to configure and manage SSPR settings:

Global Administrator

Authentication Policy Administrator

In this scenario:

Admin1 is a Global Administrator # # Allowed

Admin2 is an Authentication Policy Administrator # # Allowed

Admin3 holds Authentication Administrator and Security Administrator roles # # These roles cannot configure SSPR policies Microsoft explicitly distinguishes authentication method

management from SSPR policy configuration, which is why Authentication Administrators and Security Administrators do not have permission to configure SSPR.
Correct Users selection: Admin1 and Admin2 only
Which groups can have SSPR enabled
When enabling SSPR for Selected users, Microsoft Entra supports only assigned membership security-based groups.
According to Microsoft documentation:
Security groups (Assigned) # # Supported
Mail-enabled security groups (Assigned) # # Supported
Microsoft 365 groups # # Not supported
Dynamic groups # # Not supported
Applying this to the provided groups:
Group1 - Security, Assigned # # Eligible
Group2 - Microsoft 365, Dynamic # # Not eligible
Group3 - Mail-enabled security, Assigned # # Eligible
Correct Group selection: Group1 or Group3 only
Final Answer Summary
Section
Correct Selection
Users
Admin1 and Admin2 only
Group
Group1 or Group3 only
Microsoft Entra Administrator documentation extract:
"Only Global Administrators and Authentication Policy Administrators can configure self-service password reset settings."
"SSPR can be enabled for security groups and mail-enabled security groups with assigned membership.
Dynamic groups and Microsoft 365 groups are not supported."
Final Verified Answer:
Users: Admin1 and Admin2 only
Group: Group1 or Group3 only

NEW QUESTION: 48
You have an Azure subscription that contains the virtual networks shown in the following table.

Name	Peered with	DNS server
VNET1	VNET2	Default (Azure-provided)
VNET2	VNET1	10.10.0.4

You have the virtual machines shown in the following table.

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

You have the virtual network interfaces shown in the following table.

Name	DNS server
NIC1	Inherit from virtual network
NIC2	10.10.0.4
NIC3	Inherit from virtual network

Server1 is a DNS server that contains the resources shown in the following table.

Name	Type	Value
contoso.com	Primary DNS zone	**Not applicable**
Host1.contoso.com	A record	131.107.10.15

You have an Azure private DNS zone named contoso.com that has a virtual network link to VNET2 and the records shown in the following table.

Name	Type	Value
Host1	A record	131.107.200.20
Host2	A record	131.107.50.50

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

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: 49

Note: This question is part of a series of questions that present the same scenario. Each question

in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Microsoft Entra tenant named contoso.com.
You have a CSV file that contains the names and email addresses of 500 external users.
You need to create a guest user account in contoso.com for each of the 500 external users.
Solution; From Microsoft Entra ID in the Azure portal, you use the Bulk invite users' operation.
Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

This scenario involves inviting 500 external users (B2B guests) into your Microsoft Entra tenant (formerly Azure Active Directory). The goal is to create guest user accounts efficiently using a CSV file containing user details (names and email addresses).

Understanding the Requirement

You must onboard external users (not internal directory users). These external accounts require B2B guest invitations, not normal user creation.

There are two supported approaches for bulk guest invitations:

Using the Microsoft Entra portal (Bulk invite users)

Using PowerShell with the New-AzureADMSInvitation cmdlet.

The Bulk invite users feature in the Entra portal is specifically designed for this exact use case - bulk creation of guest accounts using a CSV file.

How the Bulk Invite Process Works

When you go to Microsoft Entra ID # Users # Bulk operations # Bulk invite, you can:

Upload a CSV file containing columns like Email address, Display name, and First/Last names.

The portal validates the file format and initiates invitations for all 500 external users.

Each external user receives an invitation email and is added as a guest (UserType = Guest) in the directory.

Once accepted, these users can be assigned roles, access resources, or be part of Azure AD groups and applications.

Why This Meets the Goal

The solution correctly leverages the Bulk invite users functionality designed for mass external user onboarding.

It satisfies all the requirements:

Creates guest accounts (not internal users).

Uses a CSV file, allowing batch import.

Achieves the goal using Microsoft Entra ID portal, without needing PowerShell scripting.

Reference from Microsoft Documentation (Microsoft Learn - Verified Source)

"You can use the Azure portal to invite multiple external users to your organization in one go using the Bulk invite feature. This process imports a CSV file containing user details and sends

invitations to all users automatically." (Source: Microsoft Learn - Bulk invite B2B users to Microsoft Entra ID)

NEW QUESTION: 50
You need to configure the Device settings to meet the technical requirements and the user requirements.
Which two settings should you modify? To answer, select the appropriate settings in the answer area.

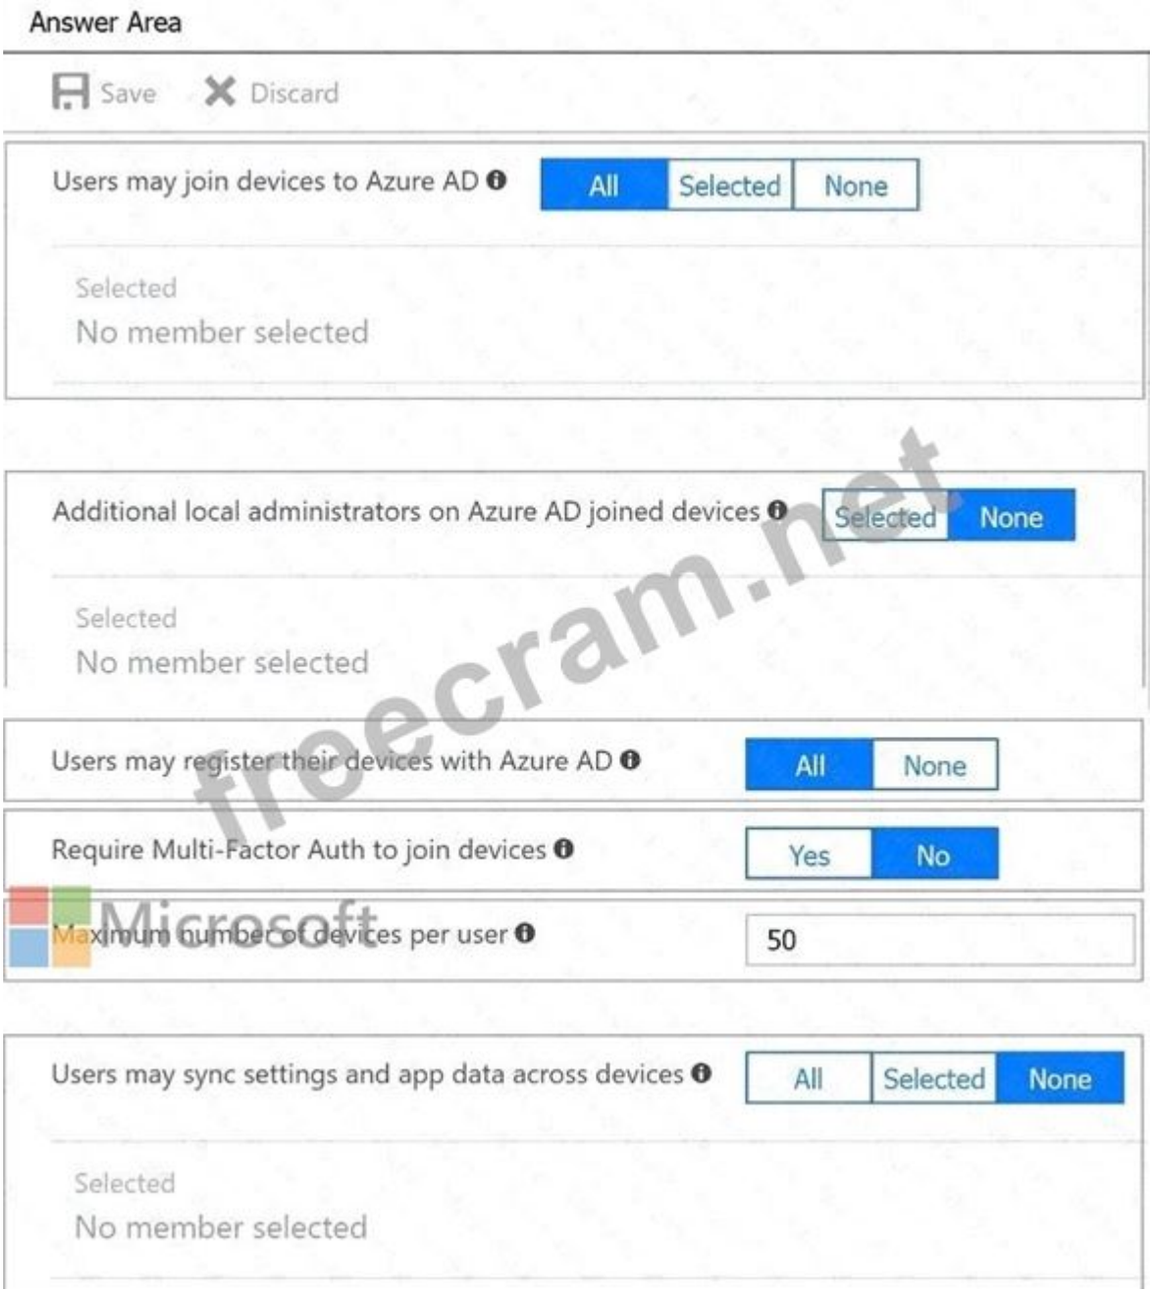

Answer:

Answer Area

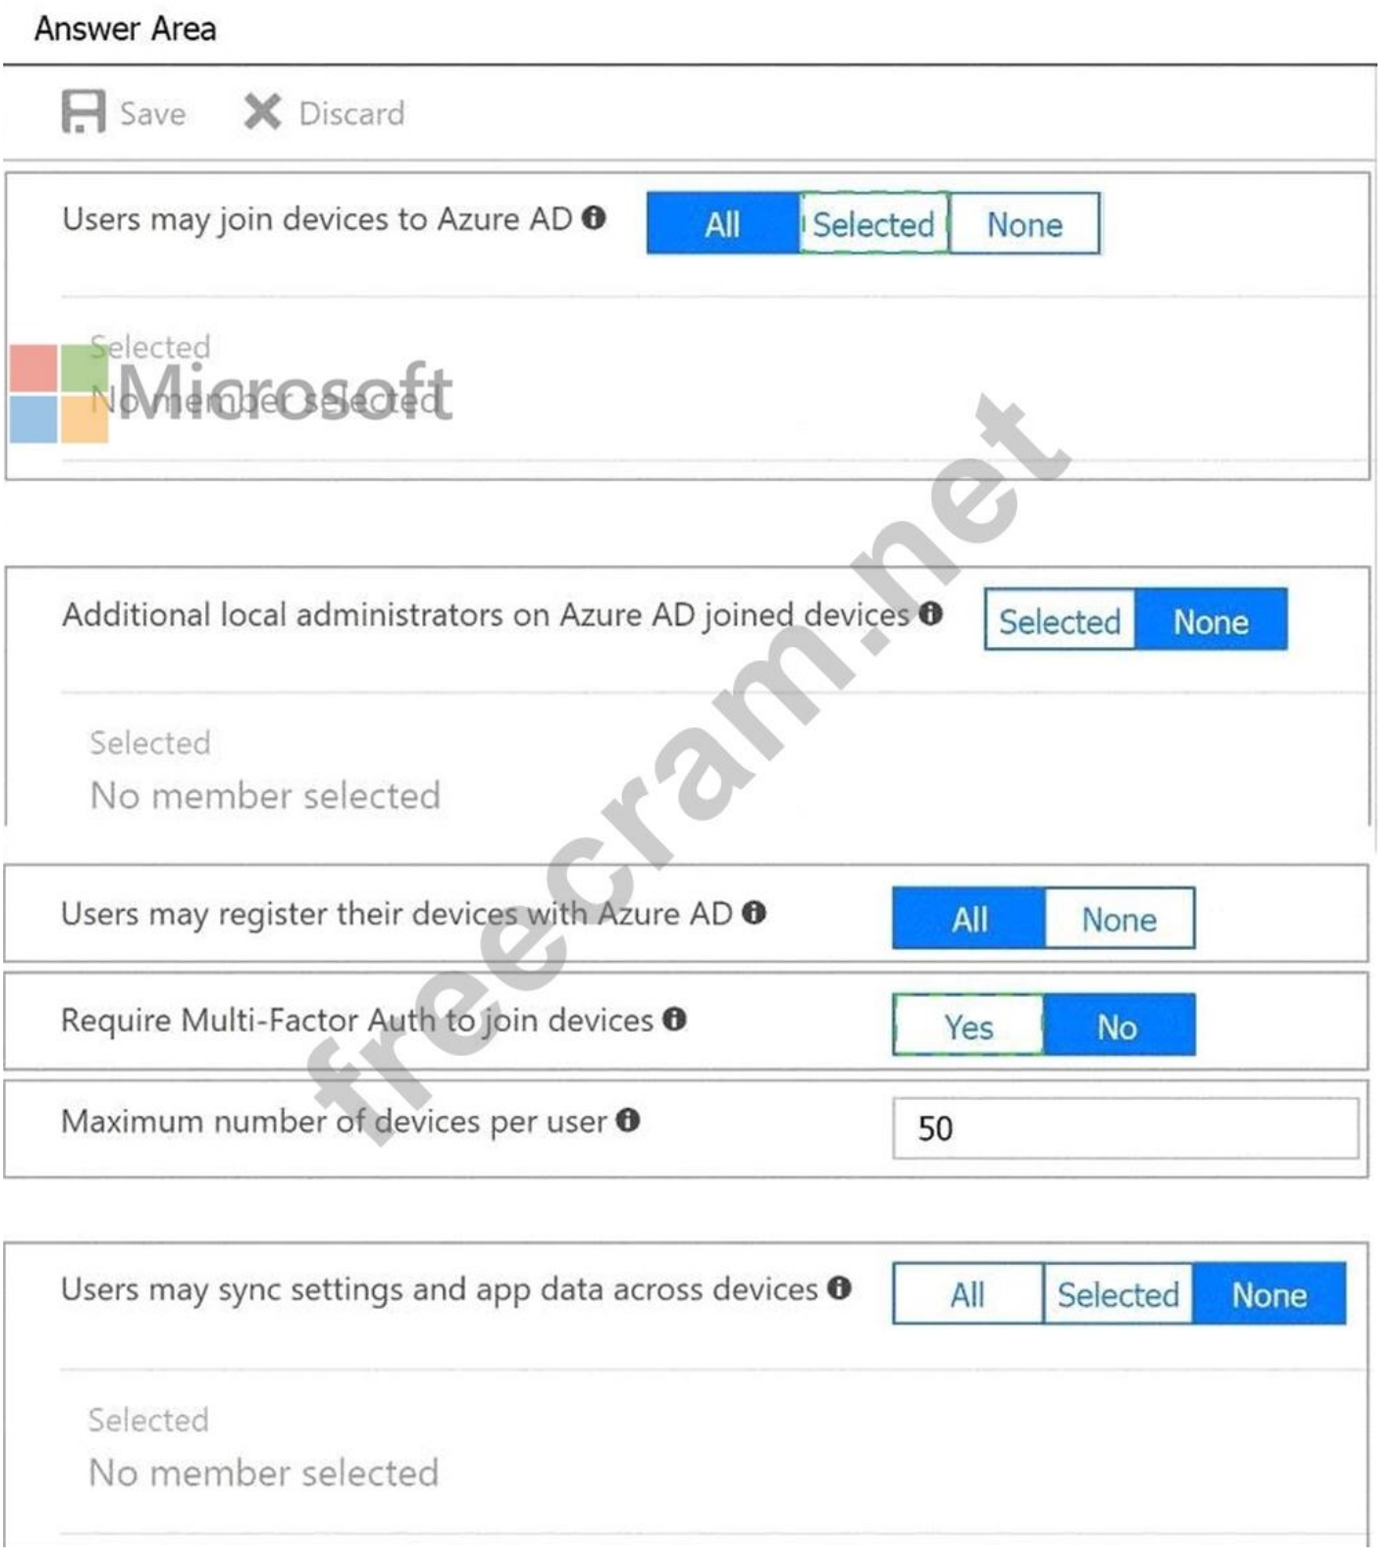

Explanation:

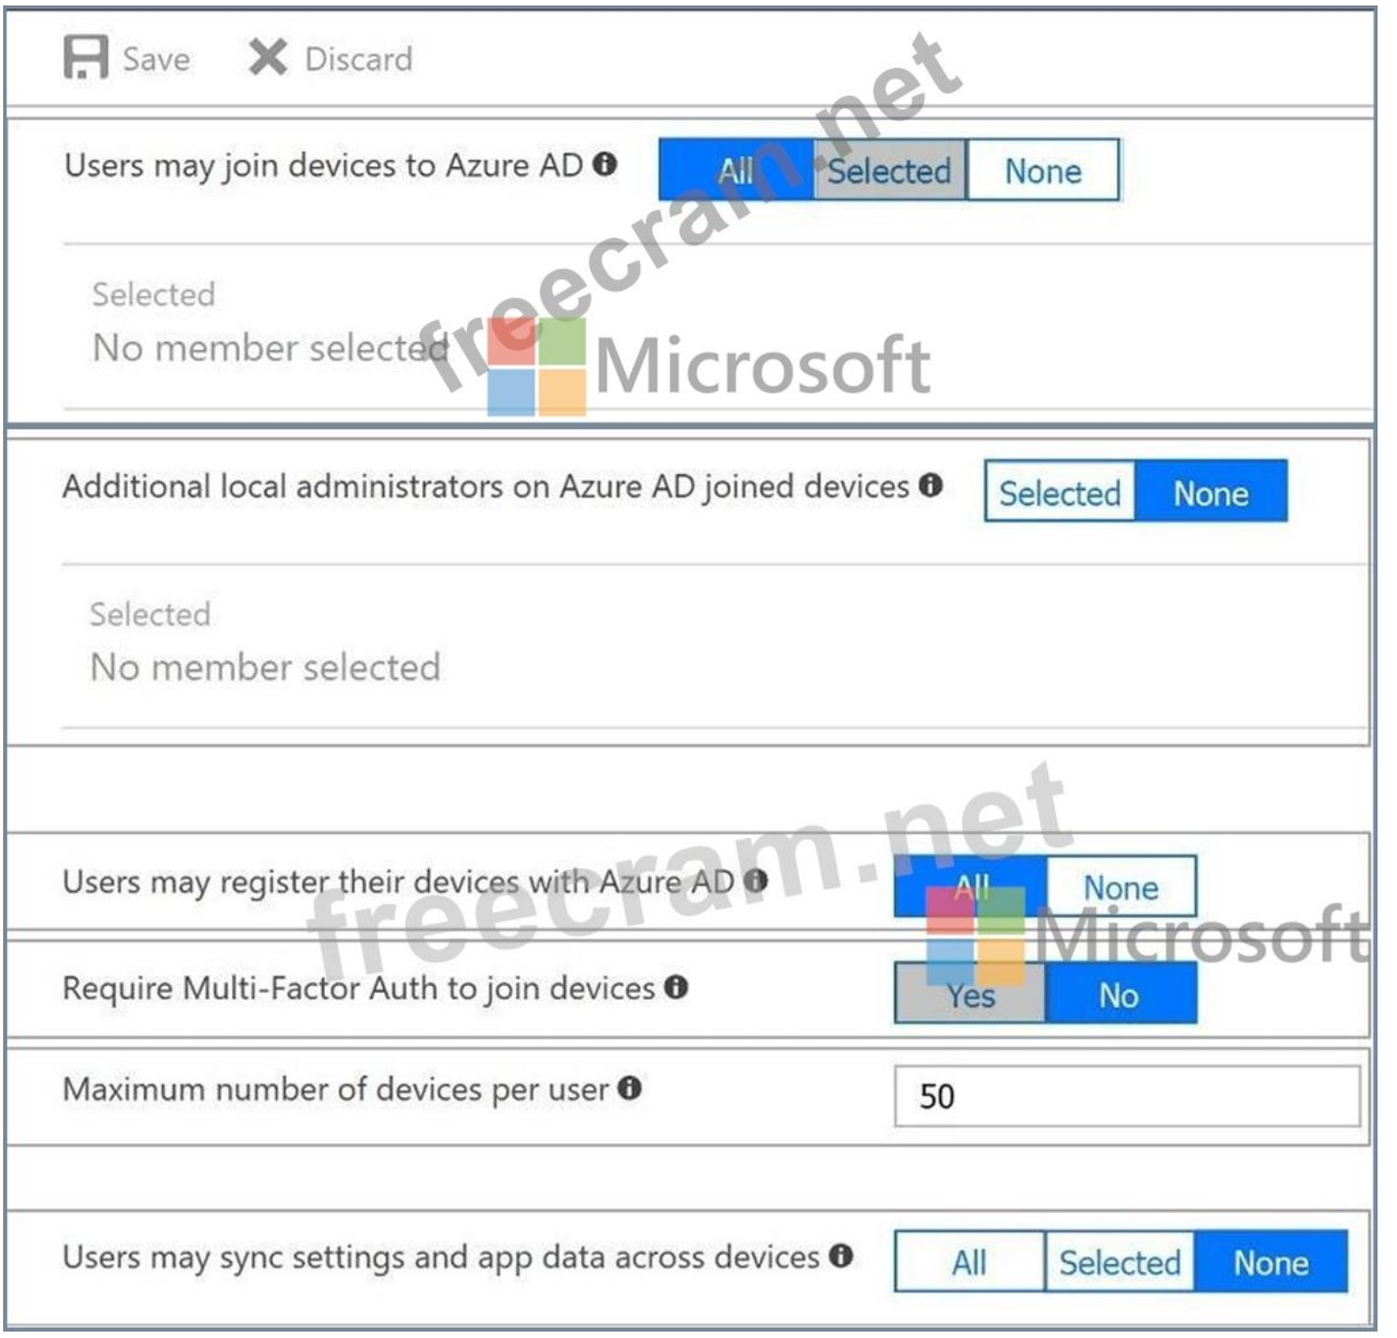

Box 1: Selected

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: 51

You have a pay-as-you-go Azure subscription that contains the virtual machines shown in the following table.

Name	Resource group	Daily cost
VM1	RG5	20 euros
VM2	RG6	30 euros

You create the budget shown in the following exhibit.

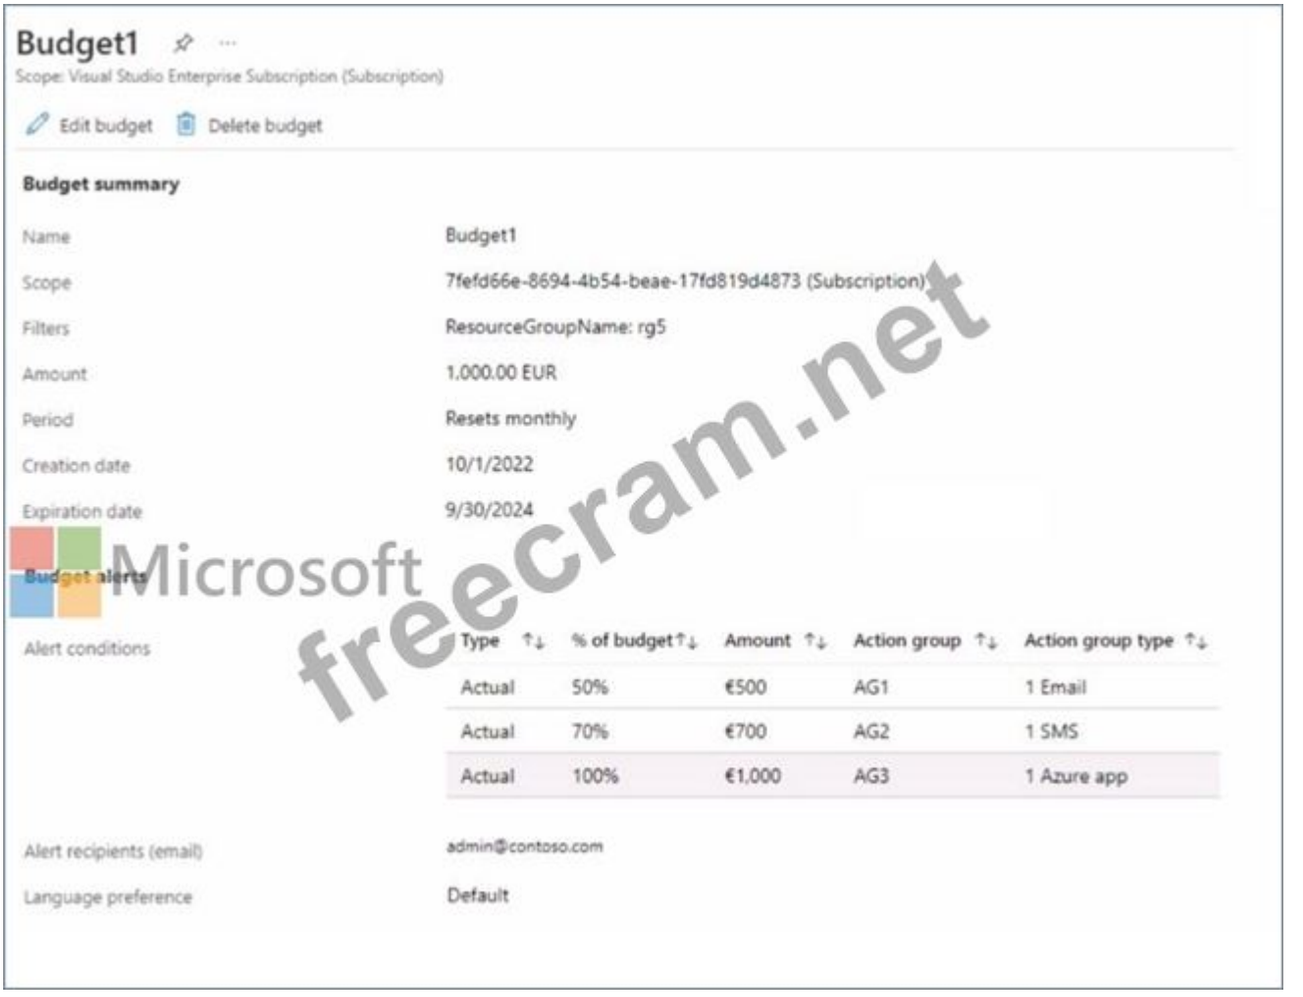

The AG1 action group contains a user named admin@contoso.com only.

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

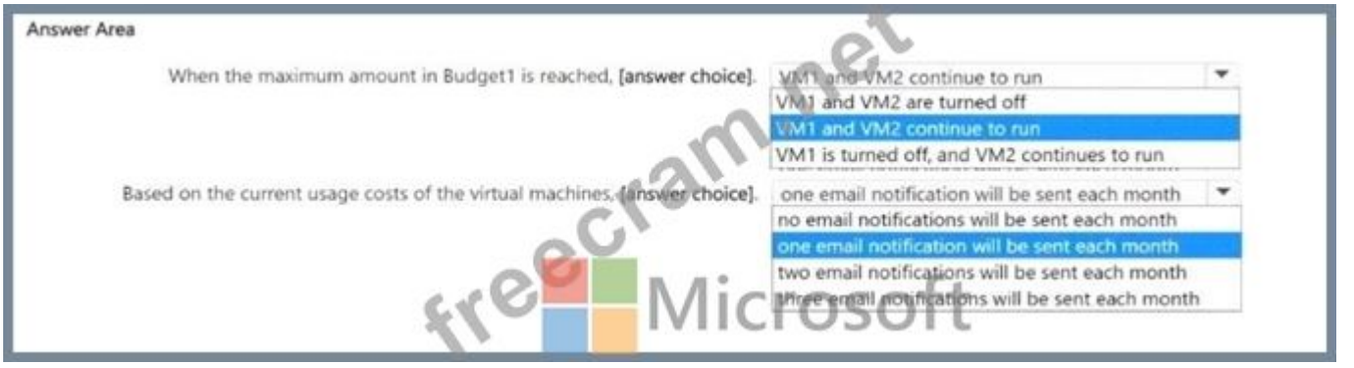

Answer:

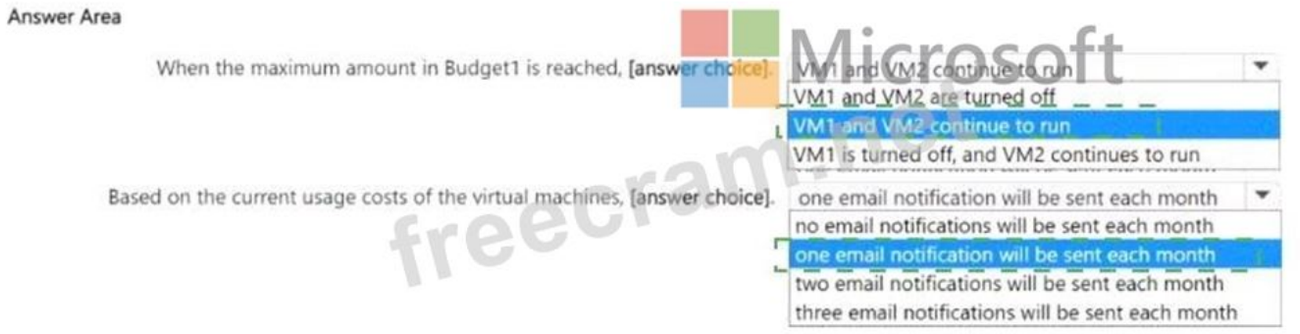

Explanation:

NEW QUESTION: 52

You have an Azure Storage account named storage1.

For storage 1. you create an encryption scope named Scope1.

Which storage types can you encrypt by using Scope1?

A. file shares only

B. containers only

C. file shares and containers only

D. containers and tables only

E. file shares, containers, and tables only

F. file shares, containers, tables, and queues

Answer: (SHOW ANSWER)

"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: 53

You have an Azure subscription. You plan to deploy the resources shown in the following table.

Name	Type
IP1	Microsoft.Network/publicIPAddresses
NSG1	Microsoft.Network/networkSecurityGroups
VNET1	Microsoft.Network/virtualNetworks
NIC1	Microsoft.Network/networkInterfaces
VM1	Microsoft.Compute/virtualMachines

You need to create a single Azure Resource Manager (ARM) template that will be used to deploy the resources. Which resource should be added to the dependsOn section for VM1?

A. IP1

B. VNET1
C. NIC1
D. NSG1
Answer: (SHOW ANSWER)
When deploying a virtual machine via an ARM template, dependencies must ensure resources are created in the proper order. A VM requires a network interface (NIC) at creation time; therefore, the VM resource must depend on the NIC.
The NIC itself depends on the virtual network, subnet, network security group, and public IP (if applicable).
Microsoft ARM template guidance specifies that the VM's dependsOn should reference the NIC, allowing Azure Resource Manager to implicitly handle the remaining network dependencies.

NEW QUESTION: 54
You have an Azure web app named App1. App1 has the deployment slots shown in the following table:

Name	Function
webapp1-prod	Production
webapp1-test	Staging

In webapp1-test, you test several changes to App1.
You back up App1.
You swap webapp1-test for webapp1-prod and discover that App1 is experiencing performance issues.
You need to revert to the previous version of App1 as quickly as possible.
What should you do?
A. Redeploy App1
B. Swap the slots
C. Clone App1
D. Restore the backup of App1
Answer: (SHOW ANSWER)
When you swap deployment slots, Azure swaps the Virtual IP addresses of the source and destination slots, thereby swapping the URLs of the slots. We can easily revert the deployment by swapping back. Deployment slots are live apps with their own host names. App content and configurations elements can be swapped between two deployment slots, including the production slot. Deploying your application to a non-production slot has the following benefits: 1. You can validate app changes in a staging deployment slot before swapping it with the production slot. 2. Deploying an app to a slot first and swapping it into production makes sure that all instances of the slot are warmed up before being swapped into production. Reference: https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots
NEW QUESTION: 55
You have an Azure subscription that contains a virtual machine named VM1 and an Azure key

vault named KV1.
You need to configure encryption for VM1. The solution must meet the following requirements:
* Store and use the encryption key in KV1.
* Maintain encryption if VM1 is downloaded from Azure.
* Encrypt both the operating system disk and the data disks.
Which encryption method should you use?
A. encryption at host
B. customer-managed keys
C. Azure Disk Encryption
D. Confidential disk encryption
Answer: (SHOW ANSWER)
Azure Disk Encryption is a service that helps you encrypt your Windows and Linux IaaS virtual machine disks1. It uses BitLocker for Windows and DM-Crypt for Linux to provide volume encryption for the OS and data disks2. Azure Disk Encryption requires that you use a key encryption key in Azure Key Vault to encrypt the volume encryption key, which is then stored on the disk. You can use either a service-managed key or a customer-managed key in Azure Key Vault3. Azure Disk Encryption also supports encrypting virtual machine disks that are downloaded from Azure4.

NEW QUESTION: 56
Peering for VNET2 is configured as shown in the following exhibit.

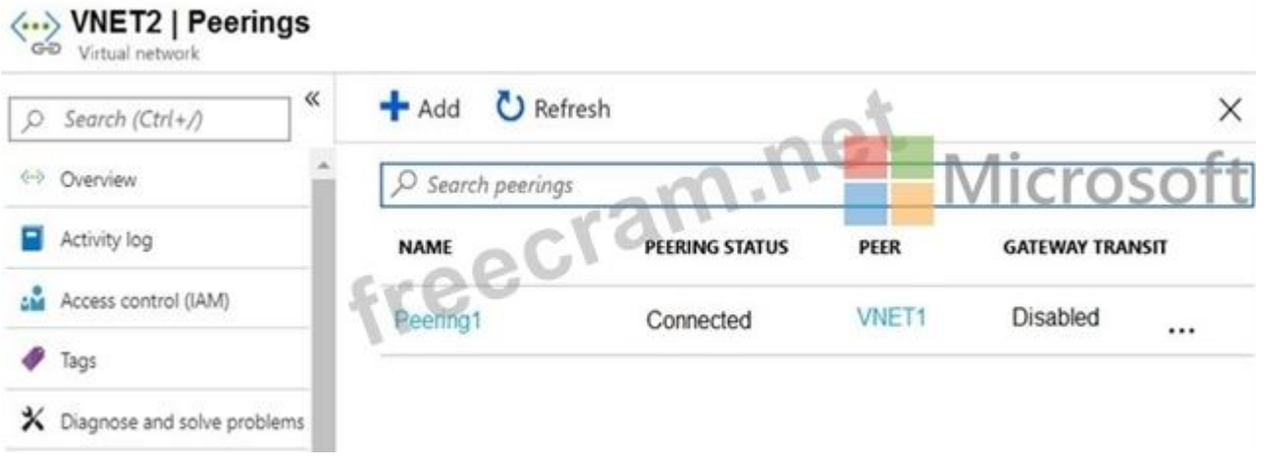

Peering for VNET3 is configured as shown in the following exhibit.

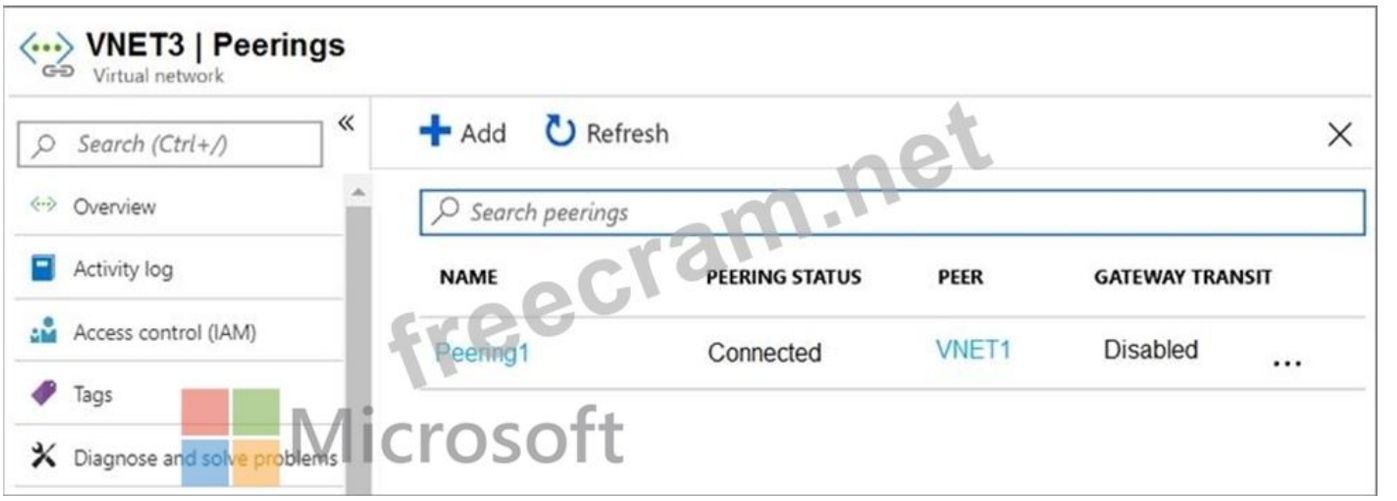

How can packets be routed between the virtual networks? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

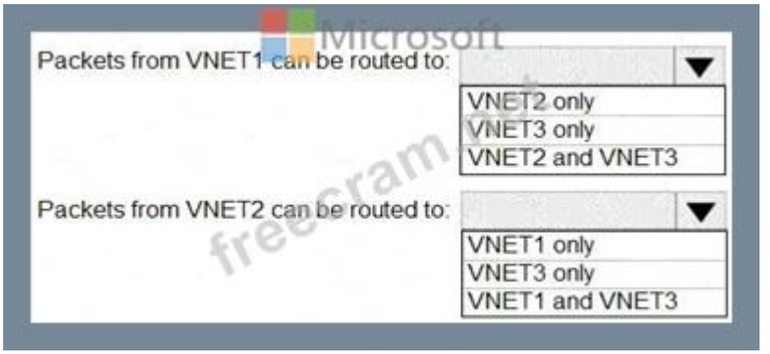

Answer:

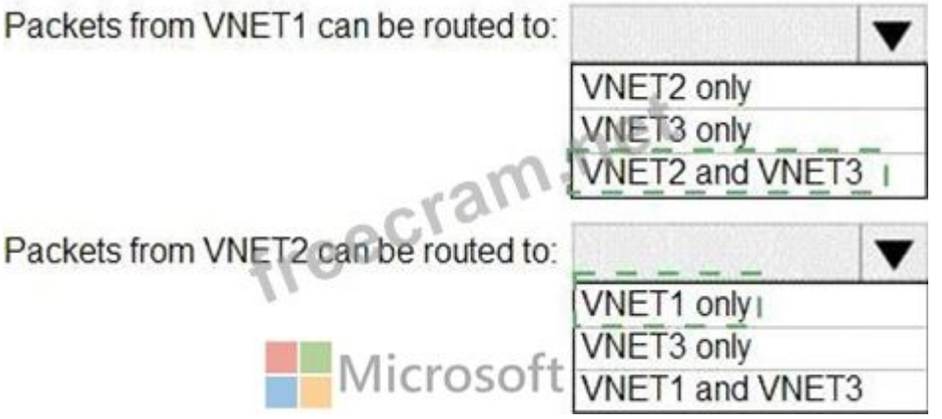

Explanation:

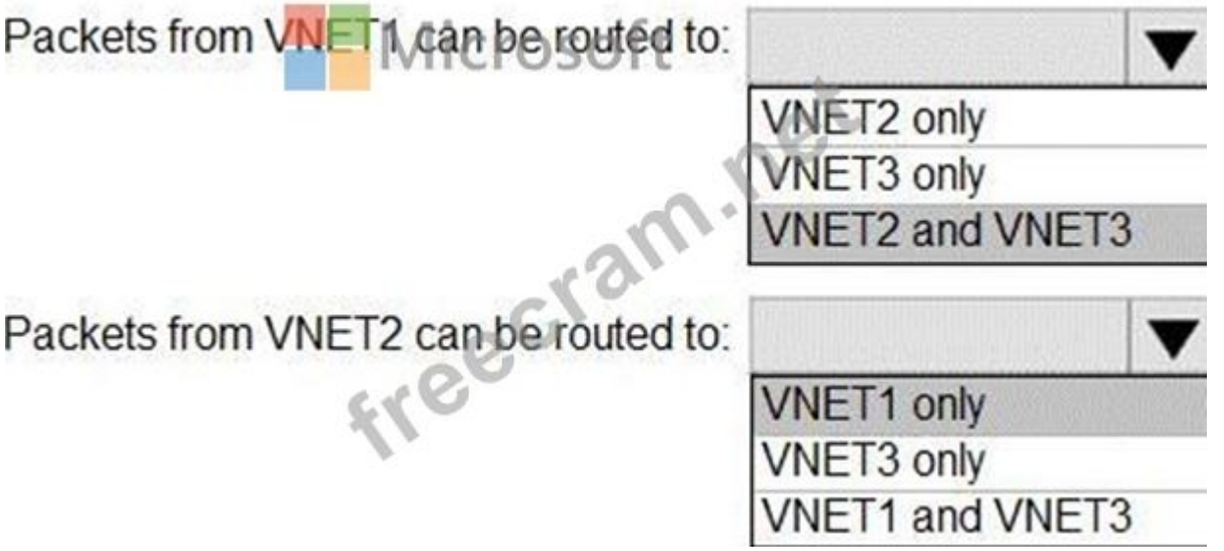

This question tests your understanding of Azure Virtual Network (VNet) Peering and its transitivity limitations.

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: 57
You have an Azure subscription named Subscription1 that has a subscription ID of c276fc76-9cd4-44c9-99a7-

4fd71546436e.

You need to create a custom RBAC role named CR1 that meets the following requirements:

Can be assigned only to the resource groups in Subscription1

Prevents the management of the access permissions for the resource groups Allows the viewing, creating, modifying, and deleting of resource within the resource groups What should you specify in the assignable scopes and the permission elements of the definition of CR1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

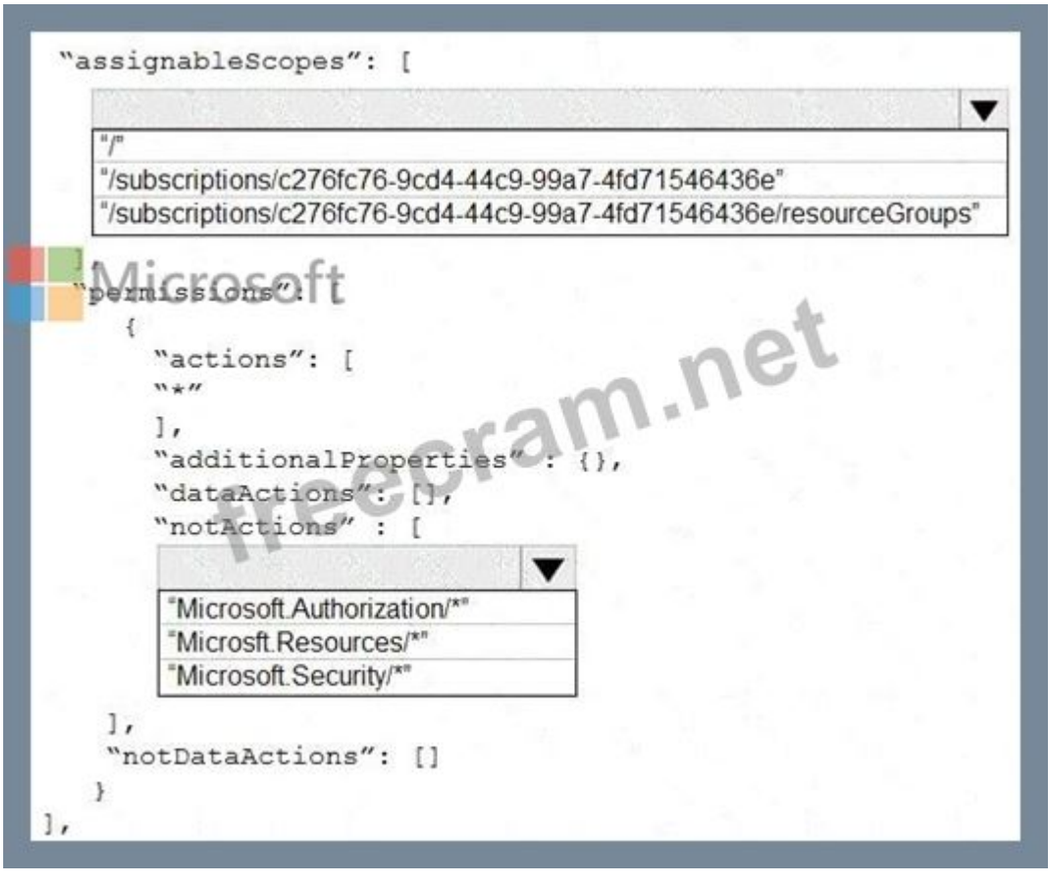

Answer:

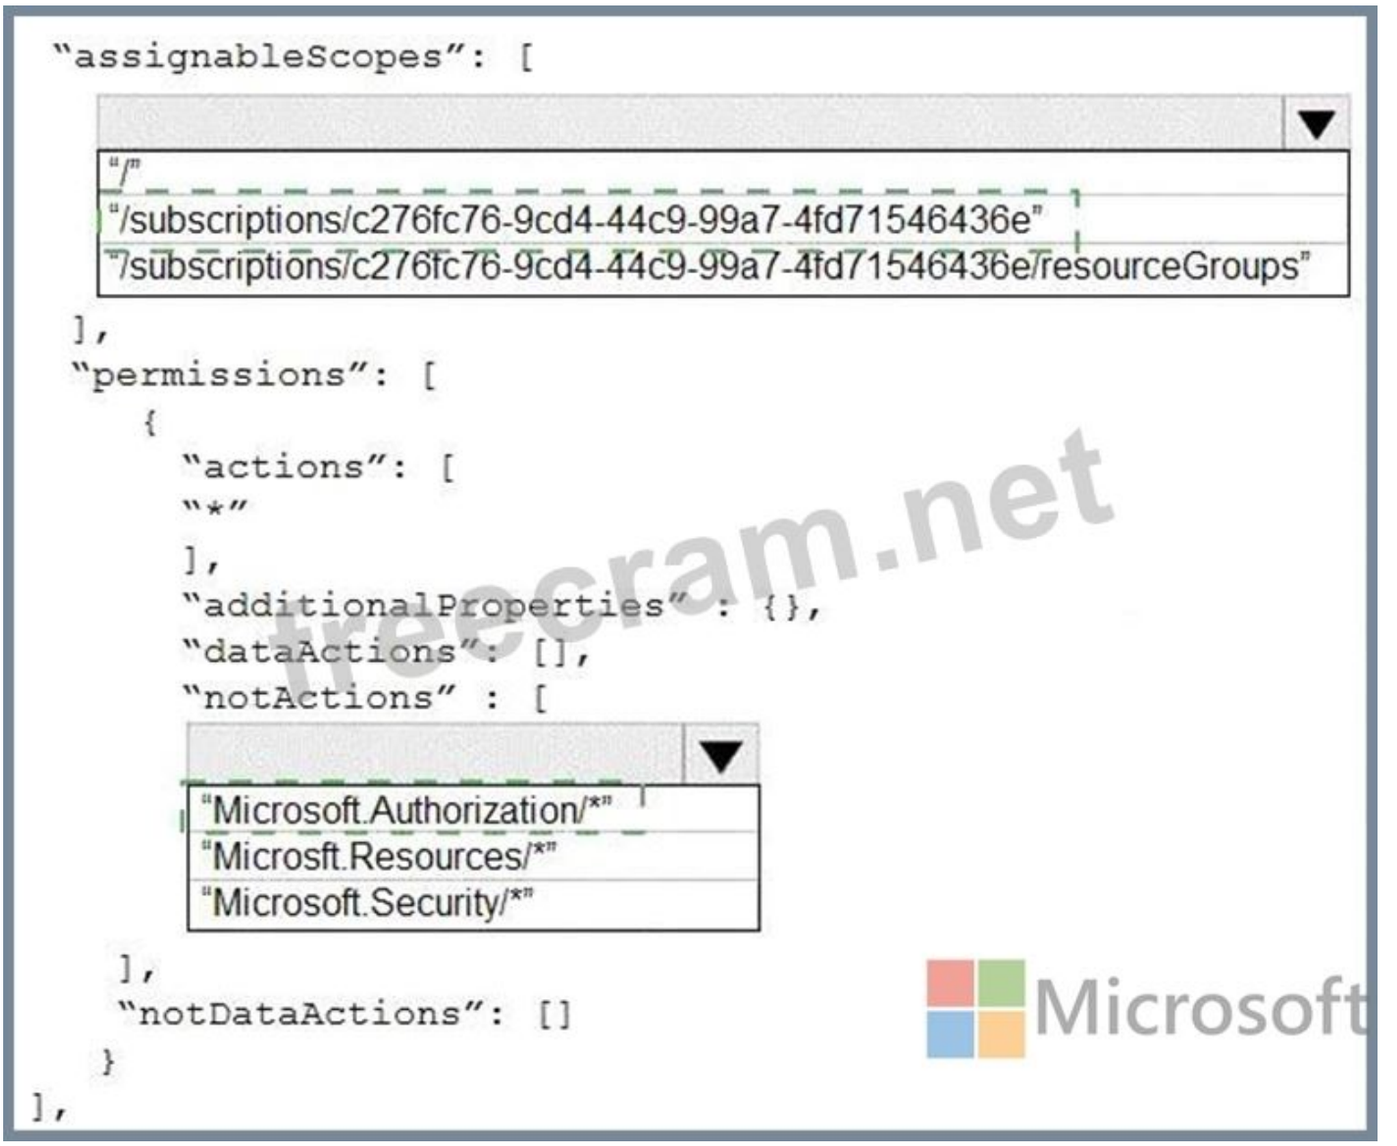

Explanation:

Box 1: "/subscription/c276fc76-9cd4-44c9-99a7-4fd71546436e"
Box 2: "Microsoft.Authorization/*"

Box 1: "/subscription/c276fc76-9cd4-44c9-99a7-4fd71546436e"

In the assignableScopes you need to mention the subscription ID where you want to implement the RBAC Box 2: "Microsoft.Authorization/*" Microsoft.Authorization/* is used to Manage authorization

Answer Area

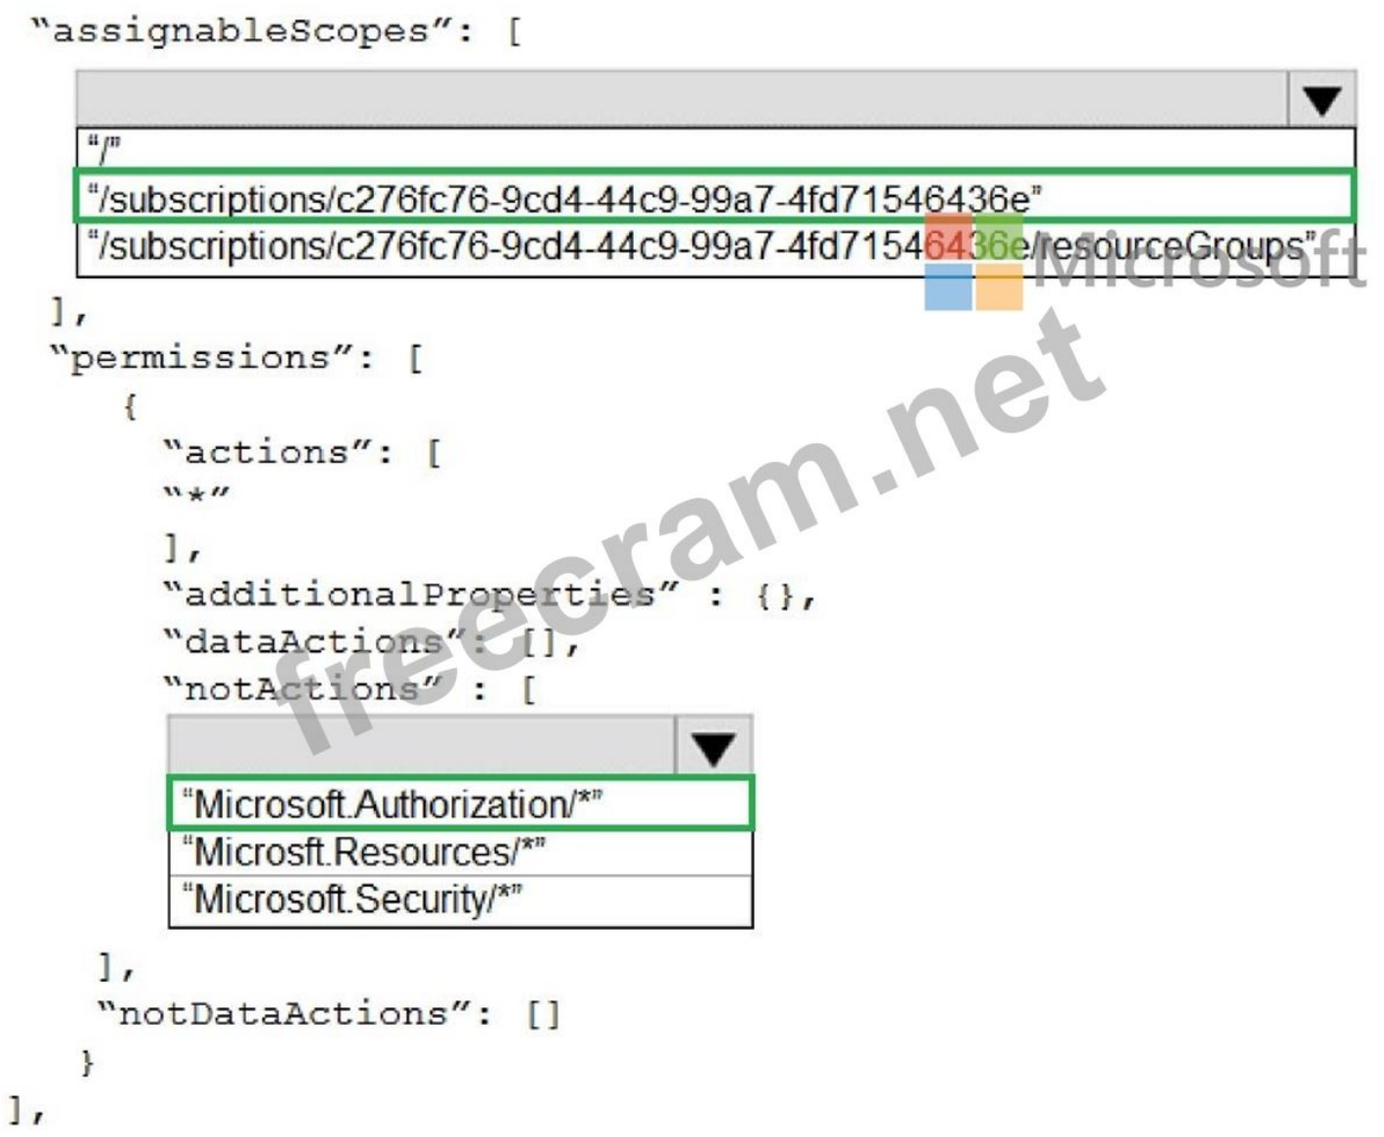

In Microsoft Azure Role-Based Access Control (RBAC), a custom role allows fine-grained permissions tailored to specific administrative needs. Custom RBAC roles are defined in JSON format and include three main elements - assignableScopes, permissions, and description.

1. Assignable Scopes

The assignable scope defines where the custom role can be applied.

In this scenario, the role must only be assignable to resource groups in the given subscription (not at the subscription or tenant root level).

Therefore, the assignable scope should be limited to the resource groups path:

"/subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e/resourceGroups"

This ensures that the role can be assigned only within resource groups in the specified subscription, not globally.

2. Permissions

The permissions section includes four arrays:

actions # defines allowed operations.

notActions # explicitly denies specific operations.

dataActions and notDataActions # for data-plane access (not needed here).
Because the requirement states that the user should be able to view, create, modify, and delete resources within the resource groups, the actions array should contain:
"actions": ["*"]
This grants full control over resource operations.
However, the question also specifies that the role must prevent the management of access permissions for the resource groups.
Managing access permissions is controlled by Microsoft.Authorization/ actions (e.g., Microsoft.Authorization
/*/Write).
Thus, to explicitly block the ability to manage permissions, the notActions array must include:
"notActions": [
"Microsoft.Authorization/*"
]
This restriction ensures that users with this role cannot assign roles, modify access control (RBAC), or change role assignments, even though they have full resource management permissions otherwise.
3. Summary of Correct JSON Configuration
{
"Name": "CR1",
"Description": "Custom role to manage resources within resource groups but not RBAC permissions",
"AssignableScopes": [
"/subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e/resourceGroups"
],
"Permissions": [
{
"Actions": [
"*"
],
"NotActions": [
"Microsoft.Authorization/*"
],
"DataActions": [],
"NotDataActions": []
}
]
}
4. Validation with Azure Administrator Study Guide
This configuration directly aligns with guidance in Microsoft Learn: "Azure custom roles in Azure RBAC" and AZ-104 Study Guide Module: Manage role-based access control (RBAC).

It meets all requirements:
Restricts role assignment scope to resource groups only.
Allows full resource management (CRUD).
Denies RBAC access control modifications via Microsoft.Authorization/*.
Final Verified Answers:
Assignable Scope: /subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e/resourceGroups
Permission (notActions): "Microsoft.Authorization/*"

NEW QUESTION: 58
You have two Azure virtual networks named VNet1 and VNet2. VNet1 contains an Azure virtual machine named VM1. VNet2 contains an Azure virtual machine named VM2.
VM1 hosts a frontend application that connects to VM2 to retrieve data.
Users report that the frontend application is slower than usual.
You need to view the average round-trip time (RTT) of the packets from VM1 to VM2.
Which Azure Network Watcher feature should you use?
A. NSG flow logs
B. Connection troubleshoot
C. IP flow verify
D. Connection monitor
Answer: (SHOW ANSWER)
https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-monitoring-overview#monitoring The connection monitor capability monitors communication at a regular interval and informs you of reachability, latency, and network topology changes between the VM and the endpoint.
Connection monitor also provides the minimum, average, and maximum latency observed over time. After learning the latency for a connection, you may find that you can decrease the latency by moving your Azure resources to different Azure regions.

NEW QUESTION: 59
You have a Microsoft Entra user named User1 and a read-access geo-redundant storage (RA-GRS) account named contoso2023. You need to meet the following requirements:
* User1 must be able to write blob data to contoso2023.
* The contoso2023 account must fail over to its secondary endpoint.
Which two settings should you configure? To answer, select the appropriate settings in the answer area.
NOTE Each correct selection is worth one point.

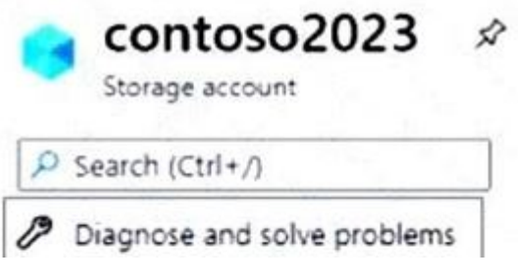


Access Control (IAM)
Data migration
Events
Storage browser
Data storage
Containers
File shares
Queues
Tables
Security + networking
Networking
Azure CDN
Access keys
Shared access signature
Encryption
Microsoft Defender for Cloud
Data management
Geo-replication
Data protection
Object replication
Blob inventory
Static website
Lifecycle management

Answer:


contoso2023
Storage account


Search (Ctrl+/)
Diagnose and solve problems
Access Control (IAM)
Data migration
Events
Storage browser
Data storage
Containers
File shares
Queues
Tables
Security + networking
Networking
Azure CDN
Access keys
Shared access signature
Encryption
Microsoft Defender for Cloud
Data management

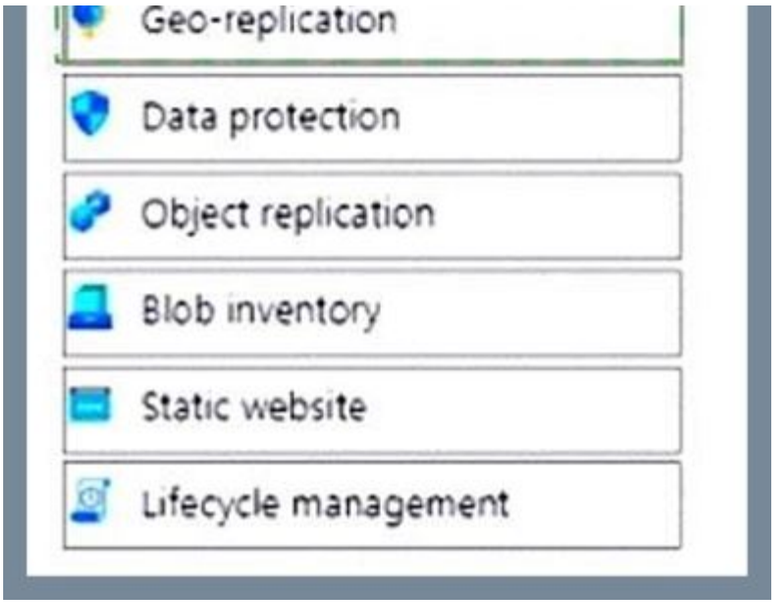

Explanation:

Access Control (IAM) - to grant User1 permission to write blob data.

Geo-replication - to perform and manage the failover operation to the secondary region.

This question involves granting write access for a user to an Azure Storage Account (RA-GRS type) and initiating a failover operation to its secondary endpoint.

Let's analyze the requirements step-by-step based on the Azure Administrator AZ-104 study guide.

Requirement 1 - "User1 must be able to write blob data to contoso2023." To allow a user (User1) to write blob data to an Azure Storage account, you must assign the appropriate Azure RBAC (Role-Based Access Control) role.

This is managed through Access Control (IAM) in the storage account.

In this scenario, you would assign a built-in Azure role such as:

Storage Blob Data Contributor (allows read/write/delete permissions to blob data) Storage Blob Data Owner (if full data control is required) This is done by navigating to:

Storage account # Access Control (IAM) # Add role assignment # Select Role # Assign to User1.

This satisfies the requirement for User1 to write blob data to the storage account.

Correct setting: Access Control (IAM)

Requirement 2 - "The contoso2023 account must fail over to its secondary endpoint." An RA-GRS (Read-Access Geo-Redundant Storage) account replicates data asynchronously to a secondary region.

Under normal circumstances, only read access is available from the secondary region.

To initiate failover, you must promote the secondary region to become the new primary.

This is achieved through the Geo-replication settings of the storage account.

Steps to perform failover:

Navigate to the storage account (contoso2023).

Under the Data management section, select Geo-replication.

Review the replication status and click Initiate account failover.

Confirm the failover operation.
This process switches the account to use the secondary endpoint as the new primary, satisfying the second requirement.
Correct setting: Geo-replication
Why Not Other Options
Access Keys / SAS: Used for authentication, not for RBAC-based permission assignment or failover operations.
Networking: Used to control access via IP or VNet, not permissions or replication.
Data Protection / Encryption: Manage backups and encryption, unrelated to failover or user access.
Verified Microsoft Learn Documentation Extract (AZ-104 Reference):
"To grant access to Azure Storage resources for specific users, assign an Azure RBAC role such as Storage Blob Data Contributor through Access Control (IAM)."
"To initiate a failover from the primary to the secondary region in a geo-redundant storage account, navigate to the Geo-replication settings and select Initiate account failover." (Source: Microsoft Learn - Manage access to Azure Storage with RBAC and Perform a storage account failover,

NEW QUESTION: 60

You have an Azure subscription that contains a virtual network named VNet1.
VNet1 uses two ExpressRoute circuits that connect to two separate on-premises datacenters.
You need to create a dashboard to display detailed metrics and a visual representation of the network topology.
What should you use?

A. Azure Monitor Network Insights
B. Azure Virtual Network Watcher
C. Log Analytics
D. a Data Collection Rule (DCR)
Answer: (SHOW ANSWER)
Azure Monitor Network Insights is specifically designed to provide end-to-end visibility into Azure networking resources, including ExpressRoute circuits, virtual networks, and on-premises connectivity. It delivers both detailed metrics and a graphical topology view, which is exactly required in this scenario.
Microsoft documentation states that Network Insights:
* Provides pre-built dashboards
* Shows visual network topology
* Supports ExpressRoute monitoring, including latency, throughput, and circuit health
* Integrates directly with Azure Monitor metrics and logs
Azure Virtual Network Watcher focuses on troubleshooting (packet capture, connection monitor) rather than dashboards. Log Analytics is query-based and lacks native topology visuals. Data Collection Rules (DCRs) define data ingestion and are not visualization tools.

NEW QUESTION: 61
You have a subnet named Subnet1 that contains Azure virtual machines. A network security group (NSG) named NSG1 is associated to Subnet1, NSG1 on default rules.
You need to create a rule in NSG1 to prevent the hosts on Subnet1 from connecting to the azure portal. The hosts must be able to connect to other ...
To what should you set Destination in the rule?
A. Service tag
B. IP addresses
C. Application security group
D. Any
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: 62
You have a Microsoft Entra tenant that contains 5,000 user accounts.
You create a new user account named AdminUser1.
You need to assign the User Administrator administrative role to AdminUser1.
What should you do from the user account properties?
A. From the Groups blade, invite the user account to a new group.
B. From the Directory role blade, modify the directory role.
C. From the Licenses blade, assign a new license.
Answer: (SHOW ANSWER)
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: 63
You have an Azure subscription that contains an Azure Backup vault named Backup1, a Recovery Services vault named Recovery1, and the resources shown in the following table.
You plan to back up the resources.
Which resource can be backed up to Backup1, and which resource can be backed up to Recovery1? To answer, select the appropriate options in the answer area.

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

NOTE: Each correct selection is worth one point.

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: 64

You have an Azure subscription linked to a hybrid Microsoft Entra tenant. The tenant contains the users shown in the following table.

Name	On-premises sync enabled
User1	No
User2	Yes

You create the Azure Files shares shown in the following table.

Name	Storage account
share1	contoso2024
share2	contoso2024
share3	contoso2025

You configure identity-based access for contoso2024 as shown in the following exhibit.

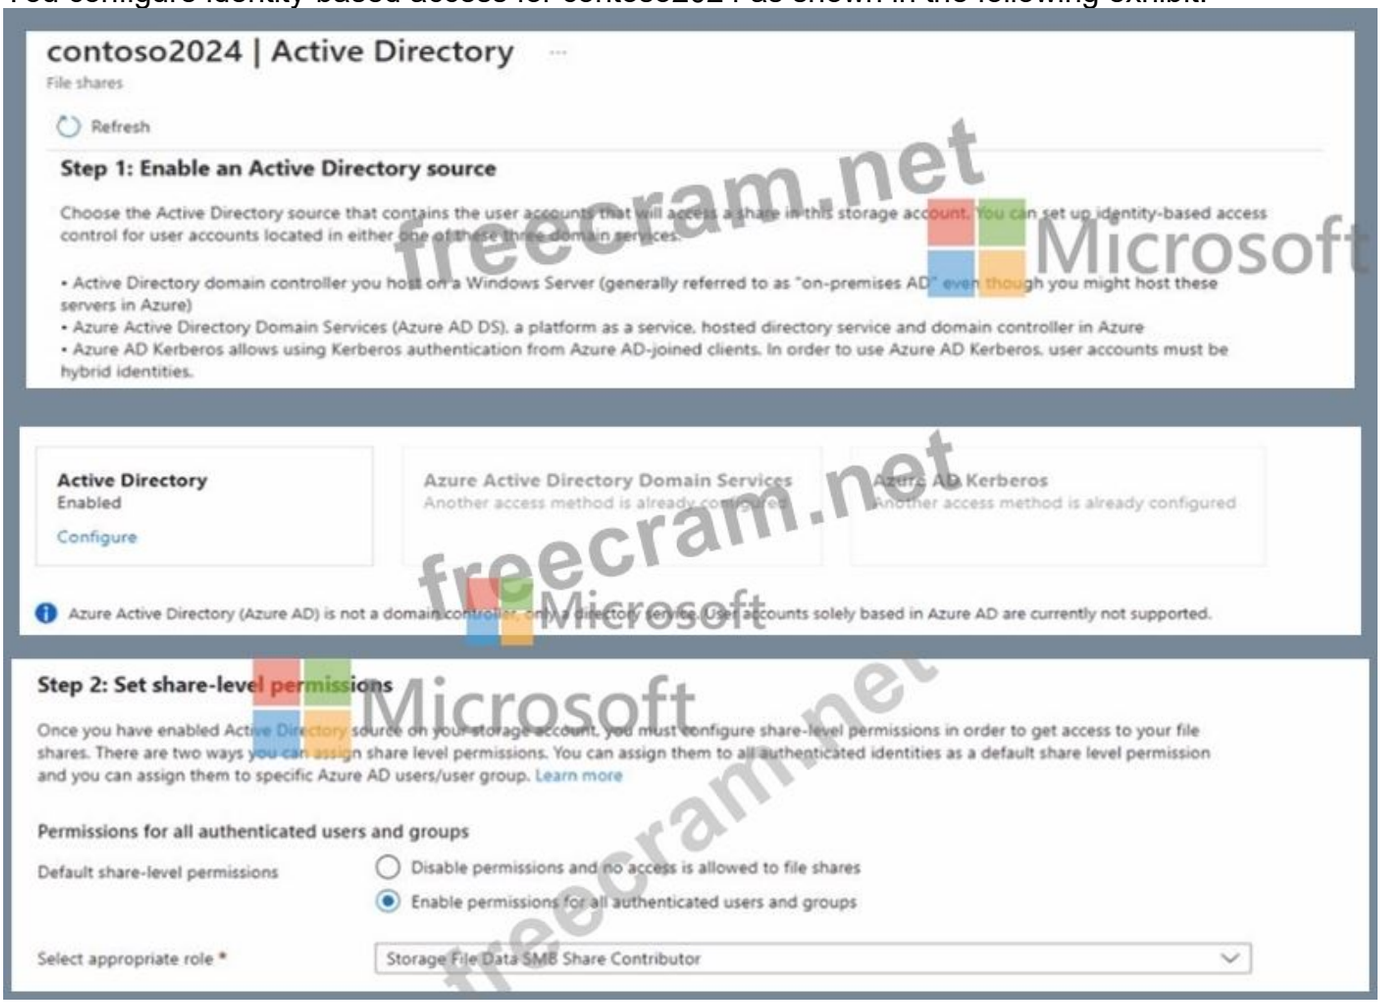

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer:

Answer Area

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: 65
You have an Azure subscription that contains the virtual networks shown in the following table.

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

The subscription contains the virtual machines shown in the following table.

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

Each virtual machine contains only a private IP address.
You create an Azure bastion for VNet1 as shown in the following exhibit.

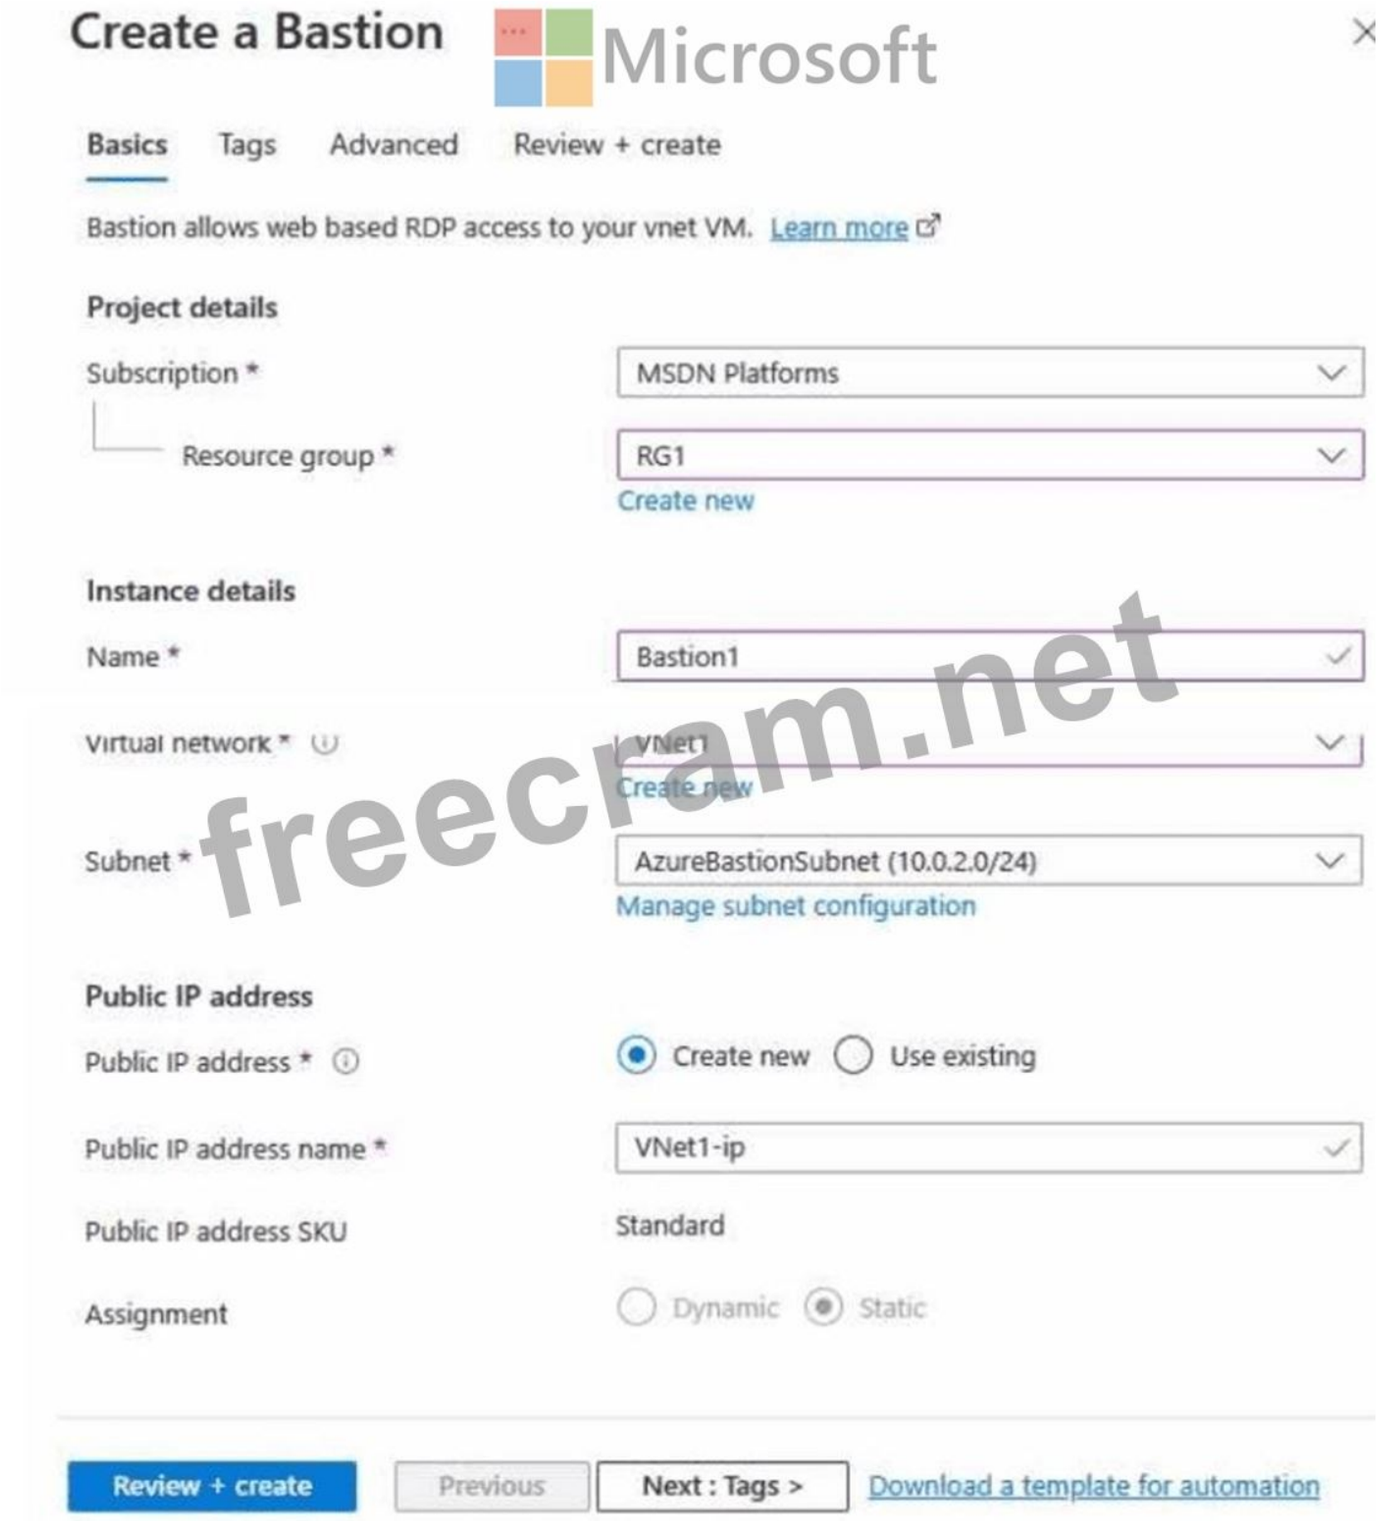

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

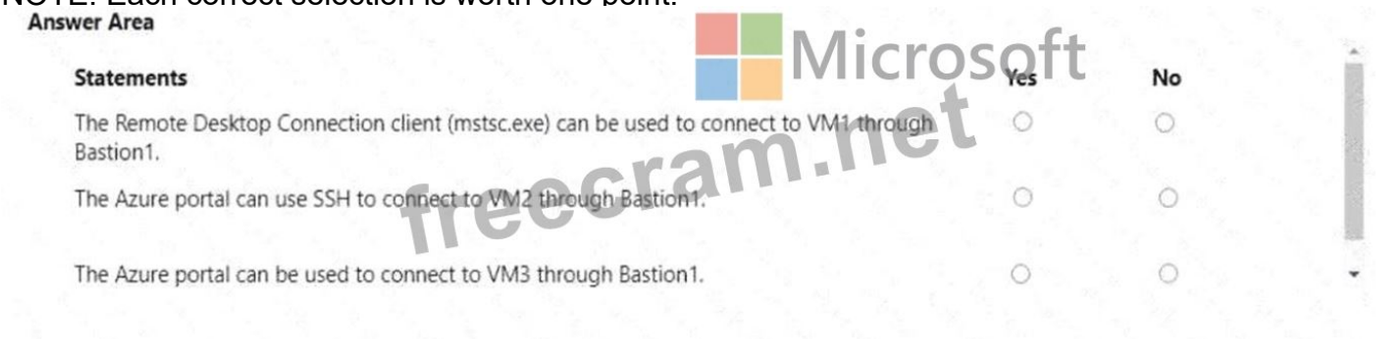

Answer:

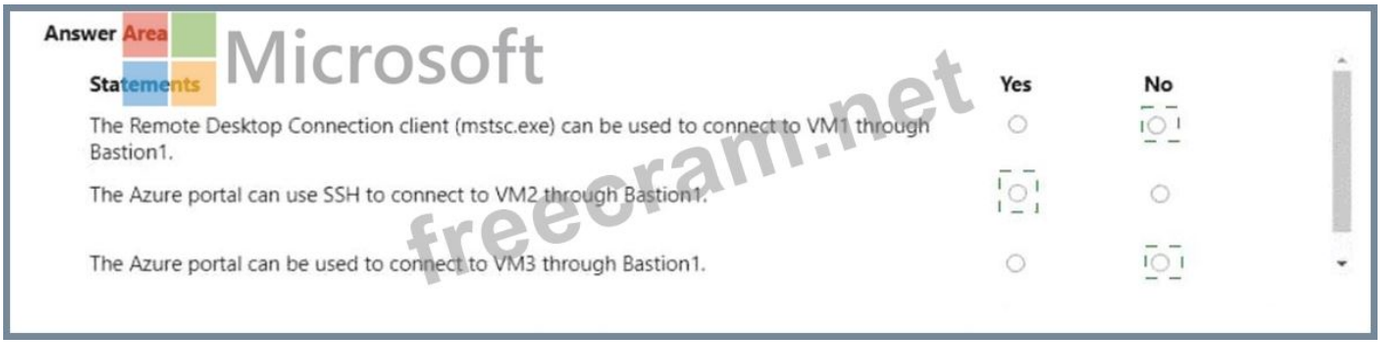

Explanation:

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: 66
You have an Azure subscription that contains the storage account shown in the following exhibit.

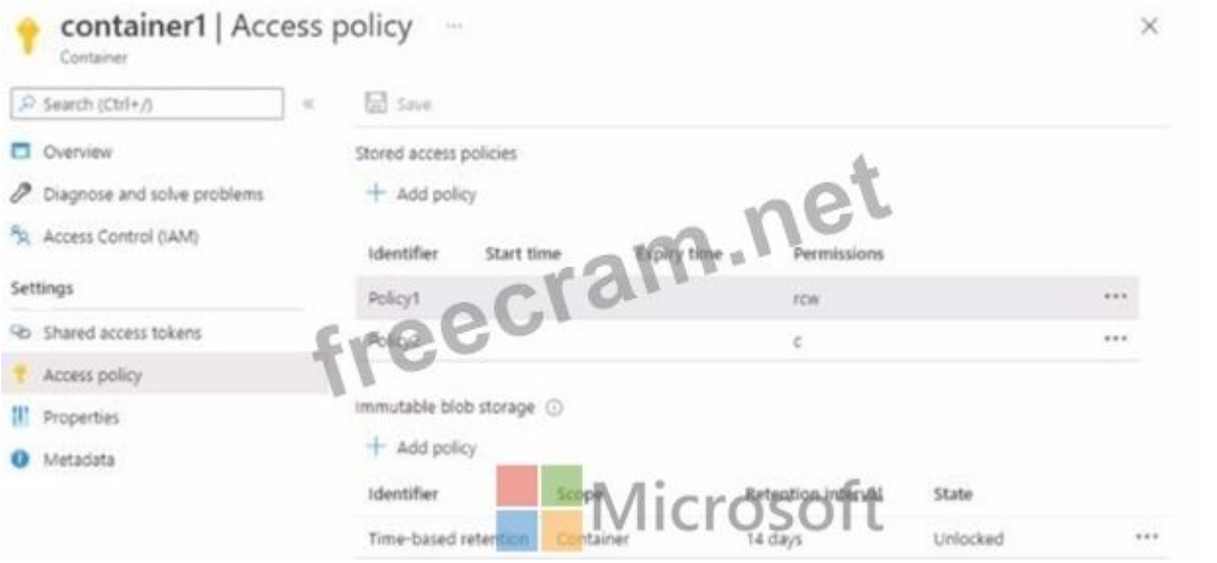

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

NEW QUESTION: 67

You have a Standard Azure App Service plan named Plan1.

You need to ensure that Plan1 will scale automatically when the CPU usage of the web app exceeds 80 percent What should you select for Plan1?

A. Automatic in the Scale out method settings

B. Rules Based m the Scale out method settings

C. Premium P1 in the Scale up (App Service plan) settings

D. Standard S1 in the Scale up (App Service plan) settings

E. Manual in the Scale out method settings

Answer: (SHOW ANSWER)

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: 68

You need to configure a new Azure App Service app named WebApp1. The solution must meet the following requirements:

* WebApp1 must be able to verify a custom domain name of app.contoso.com.

* WebApp1 must be able to automatically scale up to eight instances.

* Costs and administrative effort must be minimized.

Which pricing plan should you choose, and which type of record should you use to verify the domain? To answer, select the appropriate options in the answer area.

NOTE: Each correct answer is worth one point.

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: 69
You have an Azure Storage account named storage1 that contains two containers named container 1 and container2. Blob versioning is enabled for both containers.
You periodically take blob snapshots of critical blobs.
You create the following lifecycle management policy:

For each of the following statements, select Yes If the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

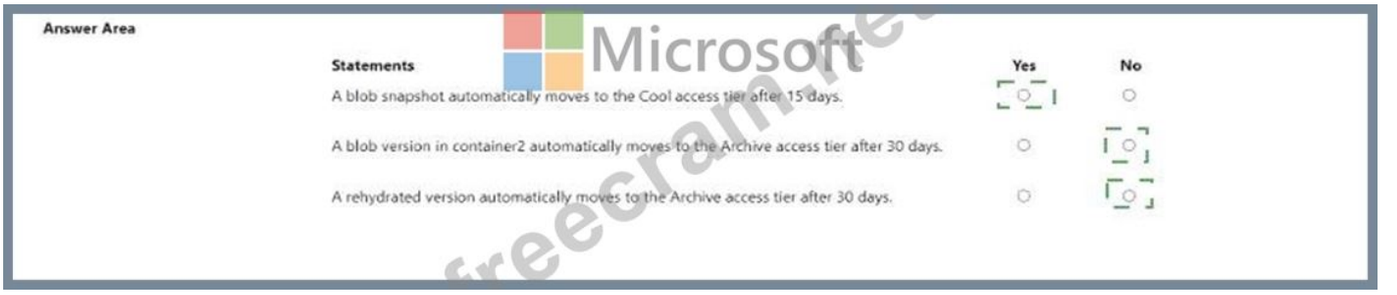

Explanation:
Based on the lifecycle management policy you created and the information from the web search results, here are the answers to your statements:
A blob snapshot automatically moves to the Cool access tier after 15 days. = Yes A blob version in container2 automatically moves to the Archive 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: 70
You have an Azure subscription that contains the resources shown in the following table:

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

You assign a policy to RG6 as shown in the following table:

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

To RG6, you apply the tag: RGroup: RG6.

You deploy a virtual network named VNET2 to RG6.

Which tags apply to VNET1 and VNET2? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

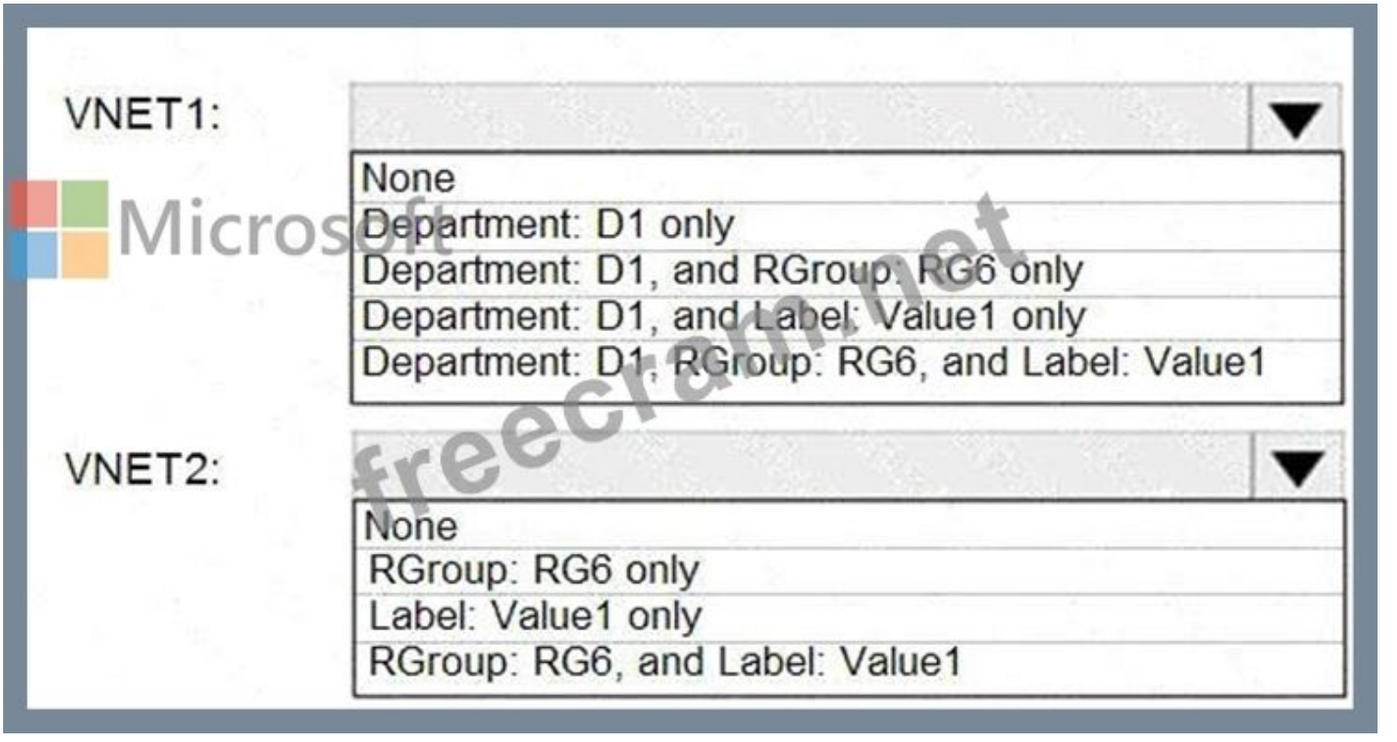

Answer:

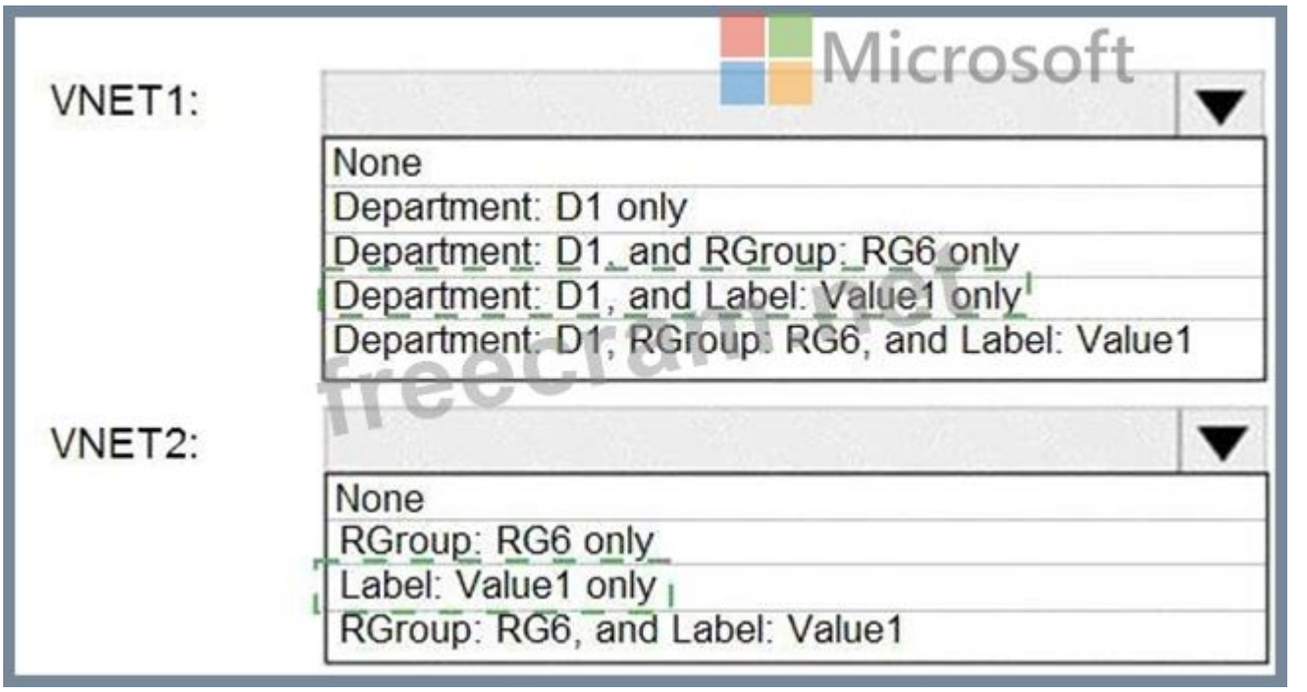

Explanation:

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: 71
You create an App Service plan named plan1 and an Azure web app named webapp1. You discover that the option to create a staging slot is unavailable. You need to create a staging slot for plan1.
What should you do first?
A. From webapp1, modify the Application settings.
B. From webapp1, add a custom domain.
C. From plan1, scale up the App Service plan.
D. From plan1, scale out the App Service plan.
Answer: (SHOW ANSWER)
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: 72
You purchase a new Azure subscription named Subscription1.
You create a virtual machine named VM1 in Subscription1. VM1 is not protected by Azure Backup.
You need to protect VM1 by using Azure Backup. Backups must be created at 01:00 and stored for 30 days.
What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

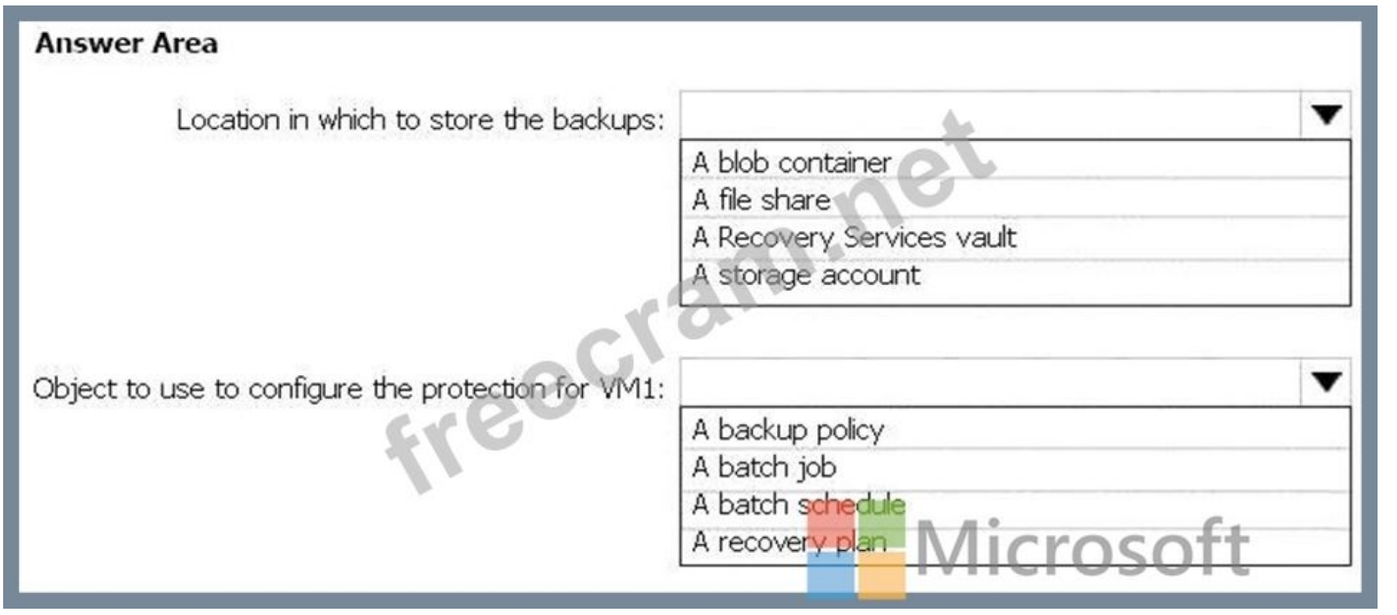

Answer:

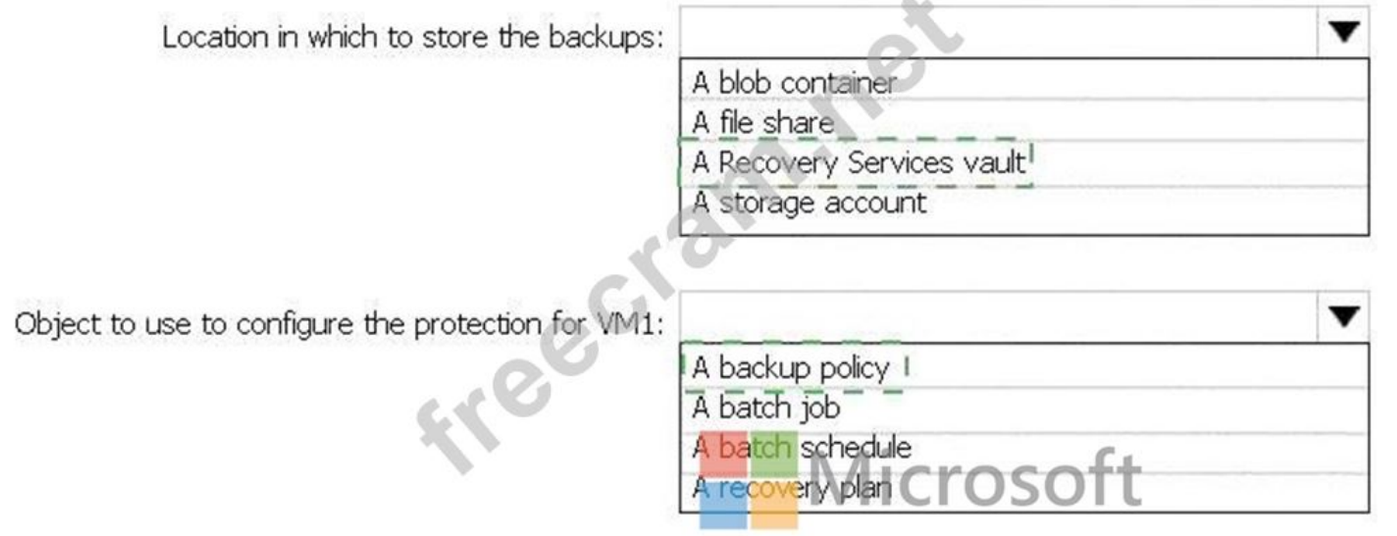

Explanation:

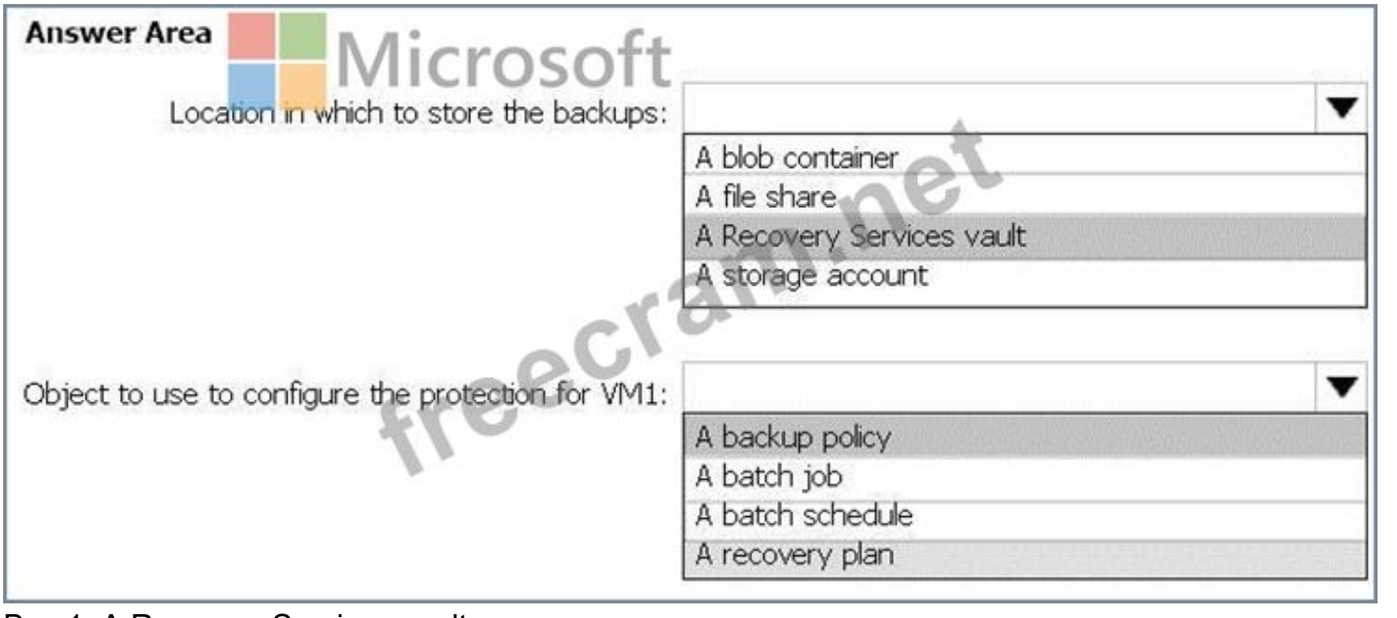

Box 1: A Recovery Services vault

A Recovery Services vault is an entity that stores all the backups and recovery points you create over time.
Box 2: A backup policy
What happens when I change my backup policy?
When a new policy is applied, schedule and retention of the new policy is followed.
References:
https://docs.microsoft.com/en-us/azure/backup/backup-configure-vault
https://docs.microsoft.com/en-us/azure/backup/backup-azure-backup-faq
A Recovery Services vault is a storage entity in Azure that houses data. The data is typically copies of data, or configuration information for virtual machines (VMs), workloads, servers, or workstations. You can use Recovery Services vaults to hold backup data for various Azure services such as IaaS VMs (Linux or Windows) and Azure SQL databases.
You can use backup policy to configure schedule.
https://docs.microsoft.com/en-us/azure/backup/backup-azure-recovery-services-vault-overviewhttps://docs.
microsoft.com/en-us/azure/backup/backup-azure-vms-first-look-arm

NEW QUESTION: 73
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure Active Directory (Azure AD) tenant named Adatum and an Azure Subscription named Subscription1. Adatum contains a group named Developers. Subscription1 contains a resource group named Dev.
You need to provide the Developers group with the ability to create Azure logic apps in the Dev resource group.
Solution: On Dev, you assign the Contributor role to the Developers group.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
The Contributor role grants the ability to create and manage all types of Azure resources, including logic apps. Assigning this role to the Developers group on the Dev resource group will allow them to create logic apps in that scope. Then, References: [Built-in roles for Azure resources] [Azure Logic Apps permissions and access control]

NEW QUESTION: 74
You have an Azure virtual network named VNetl that contains the following settings:
* IPv4 address space: 172.16.10.0/24

* Subnet name: Subnetl
* Subnet address range: 172.16.10.0/25
What is the maximum number of virtual machines that can connect to 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: 75
You have an Azure Storage account named storage1.
You have an Azure App Service app named app1 and an app named App2 that runs in an Azure container instance. Each app uses a managed identity.
You need to ensure that App1 and App2 can read blobs from storage1 for the next 30 days.
What should you configure in storage1 for each app?

Answer:

Explanation:

Box 1: Access Control (IAM)

Since the App1 uses Managed Identity, App1 can access the Storage Account via IAM. As per requirement, we need to minimize the number of secrets used, so Access keys is not ideal.

Box 2: Shared access signatures (SAS)

We need temp access for App2, so we need to use SAS.

A shared access signature (SAS) provides secure delegated access to resources in your storage account without compromising the security of your data. With a SAS, you have granular control over how a client can access your data. You can control what resources the client may access, what permissions they have on those resources, and how long the SAS is valid, among other parameters.

Reference:

https://docs.microsoft.com/en-us/azure/storage/common/storage-sas-overview

https://docs.microsoft.com/en-us/azure/storage/common/storage-auth

NEW QUESTION: 76

You have an Azure container registry in the Standard tier.

You need to ensure that you can configure Azure Private Link for the container registry. What should you do first?

A. Create a new virtual network.

B. Configure the Access keys settings.

C. Upgrade the container registry to the Premium tier.

D. Configure the Access Control (1AM) settings.

Answer: (SHOW ANSWER)

Azure Container Registry (ACR) is a managed, private Docker registry service that allows you to build, store, and manage container images and artifacts. Microsoft provides three service tiers for ACR - Basic, Standard, and Premium - each offering different capabilities and performance features.

According to the Microsoft Azure Administrator documentation, Private Link integration is only available in the Premium tier of Azure Container Registry. Private Link enables you to connect to Azure services (like ACR) through a private endpoint in your virtual network, ensuring that traffic between your virtual network and the registry remains entirely within the Microsoft backbone network. This prevents exposure to the public internet and significantly enhances the security of image pull and push operations.
In the Standard and Basic tiers, ACR only supports access through public endpoints, optionally restricted by IP firewall rules or service endpoints. To configure Azure Private Link, the registry must first be upgraded to Premium tier, after which you can create a Private Endpoint associated with your registry.
Once upgraded, you can configure network rules, DNS resolution, and endpoint connections under the ACR's networking settings to fully implement private access to the registry.
Thus, the verified and Microsoft-documented answer is to upgrade the container registry to the Premium tier before configuring Azure Private Link.

NEW QUESTION: 77
You have an Azure subscription that contains a virtual network named VNet1. VNet1 uses an IP address space of 10.0.0.0/16 and contains the subnets in the following table.

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 contains a virtual appliance named VM1 that operates as a router.
You create a routing table named RT1.
You need to route all inbound traffic to VNet1 through VM1.
How should you configure RT1? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

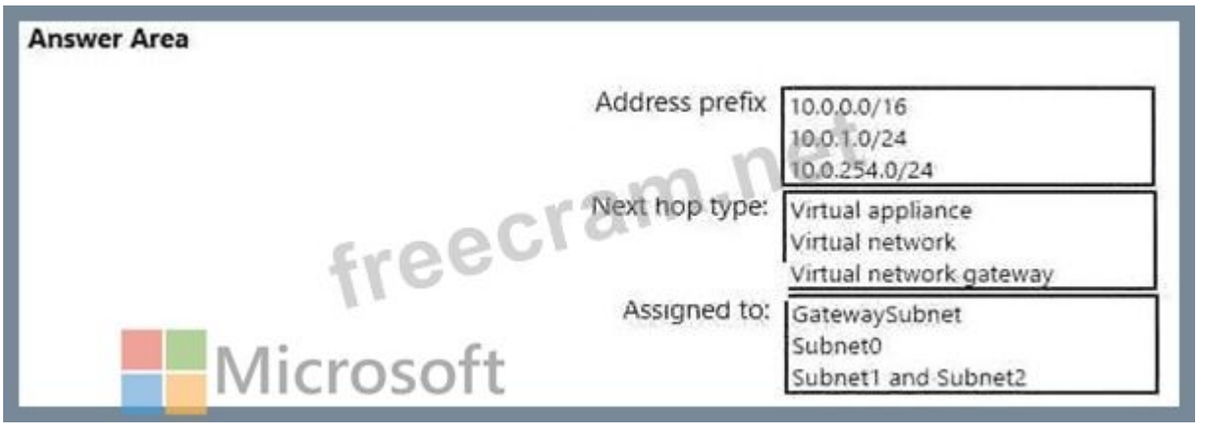

Answer:

Answer Area

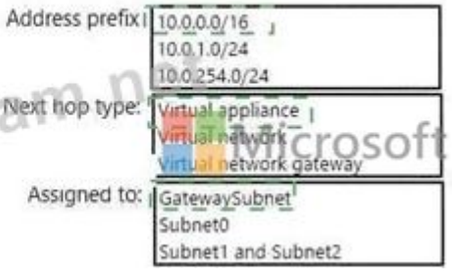

Explanation:

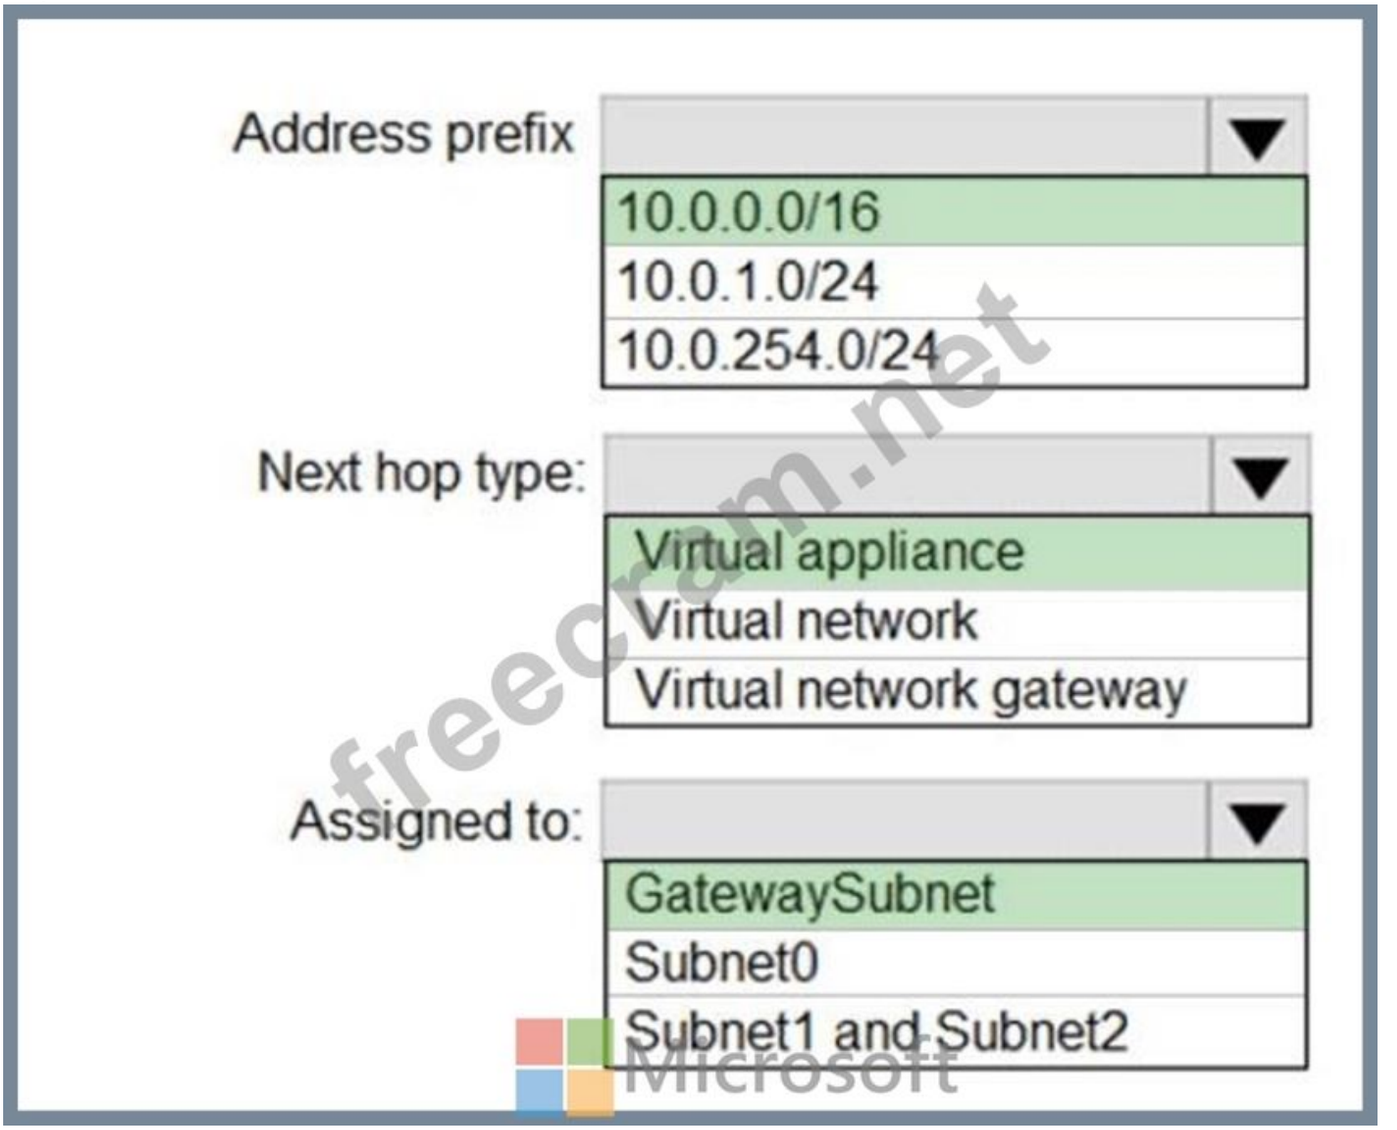

Box1 : 10.0.0.0/16
Address prefix in networking refer to the destination IP address range. In this scenario, destination is Vnet1 , hence Address prefix will be the address space of Vnet1.
Box 2 : Virtual appliance
Next hop gets the next hop type and IP address of a packet from a specific VM and NIC. 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: 78
You have a Microsoft Entra tenant that contains a user named External User External User authenticates to the tenant by using extemall95@gmail.com.
You need to ensure that External User authenticates to the tenant by using contractor@gmail.com.
Which two settings should you configure from the Overview blade? To answer, select the appropriate settings in the answer area NOTE: Each correct answer is worth one point.

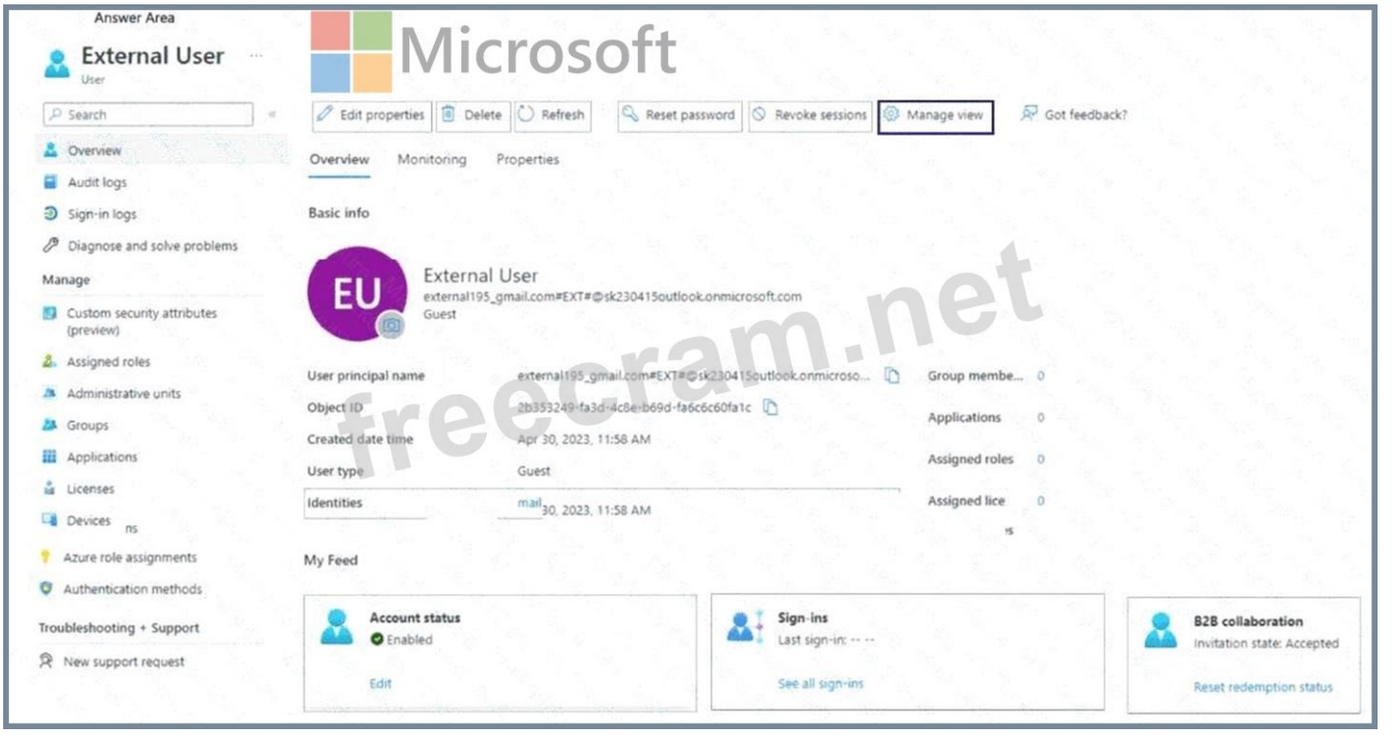

Answer:

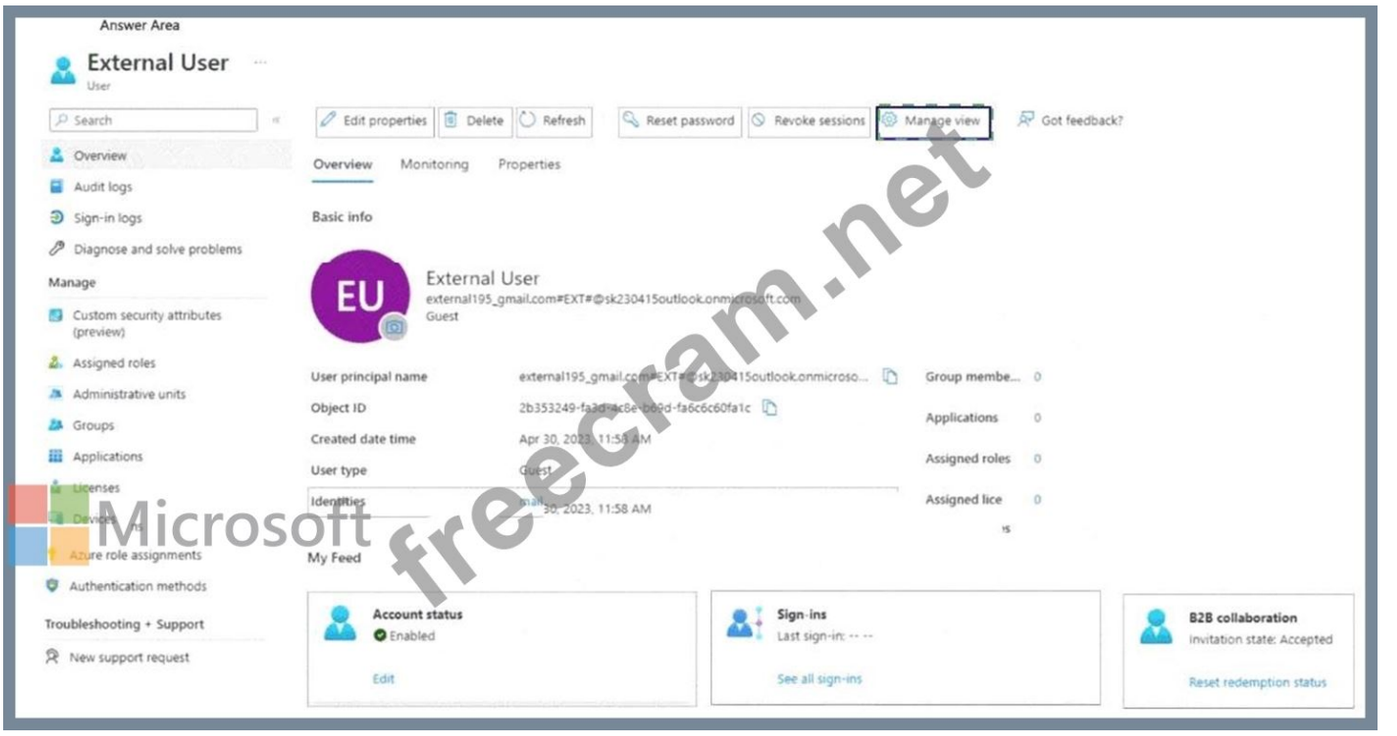

Explanation:

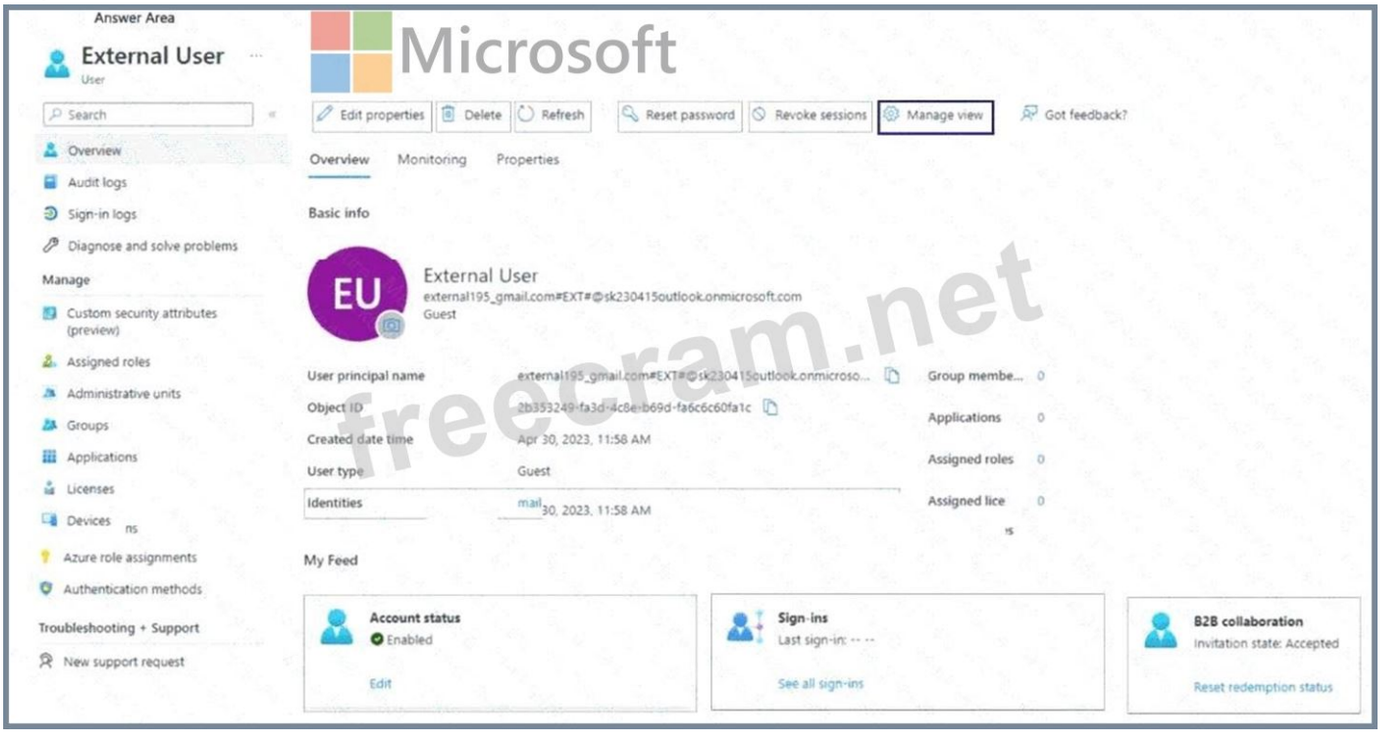

In Microsoft Entra ID (formerly Azure Active Directory), guest users (External Users) are typically added via B2B collaboration invitations. They authenticate using the identity provider (IdP) associated with the email they were invited with - such as a Microsoft account, Gmail (Google), or another organization's 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: 79
You have an Azure subscription that contains multiple virtual machines in the West US Azure region.
You need to use Traffic Analytics in Azure Network Watcher to monitor virtual machine traffic.
Which two resources should you create? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. a Data Collection Rule (OCR) in Azure Monitor
B. a Log Analytics workspace
C. an Azure Monitor workbook
D. a storage account
E. a Microsoft Sentinel workspace
Answer: (SHOW ANSWER)
Traffic Analytics is a feature within Azure Network Watcher that provides insights into network traffic patterns, security threats, and performance using NSG flow logs. To enable Traffic Analytics, two key resources are required:
* A Log Analytics Workspace - This is where Azure Network Watcher sends and stores the processed flow log data for analysis.
* A Storage Account - Network Security Group (NSG) flow logs are first written to a designated storage account before they are ingested by Traffic Analytics for processing.
According to Microsoft Azure Administrator documentation, under "Enable Traffic Analytics":
"Traffic Analytics uses Network Watcher NSG flow logs, which must be stored in a storage account. These logs are then processed and analyzed in a Log Analytics workspace to provide

insights into network activity." Workflow:
* NSG flow logs # stored in Azure Storage Account
* Logs are processed by Azure Monitor # ingested into Log Analytics workspace
* Traffic Analytics visualizes and analyzes the data
Other options such as Azure Monitor workbooks or Sentinel workspaces are not mandatory to activate Traffic Analytics.
Thus, to use Traffic Analytics, you must create:
A Storage Account to store flow logs
A Log Analytics Workspace to process and visualize them

NEW QUESTION: 80

You download an Azure Resource Manager template based on an existing virtual machine. The template will be used to deploy 100 virtual machines.
You need to modify the template to reference an administrative password. You must prevent the password from being stored in plain text.
What should you create to store the password?

A. Azure Active Directory (AD) Identity Protection and an Azure policy
B. a Recovery Services vault and a backup policy
C. an Azure Key Vault and an access policy
D. an Azure Storage account and an access policy
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: 81
You have an Azure subscription named Subscription1 that contains a virtual network named VNet1. You add the users in the following table.

User	Role
User1	Owner
User2	Security Admin
User3	Network Contributor

Which user can perform each configuration? To answer select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
Answer Area

Answer:
Answer Area

Explanation:

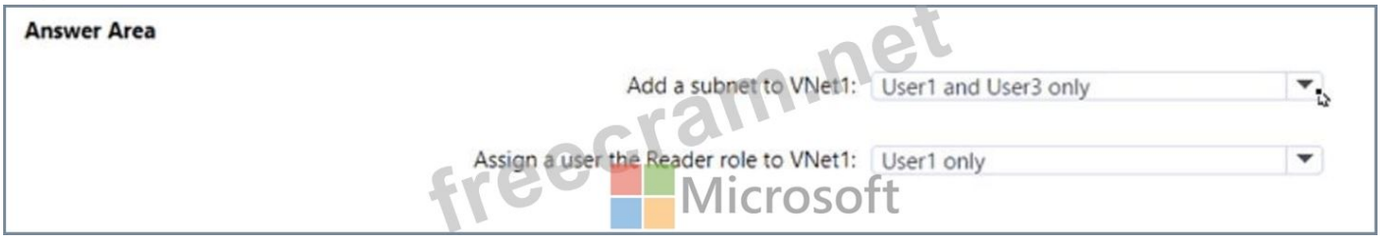

This question evaluates understanding of Azure Role-Based Access Control (RBAC) permissions for managing virtual networks (VNets) and assigning Azure roles.

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: 82
You have an Azure subscription that contains a resource group named RG1.
You plan to create an Azure Resource Manager (ARM) template to deploy a new virtual machine named VM1. VM1 must support the capture of performance data.
You need to specify resource dependencies for the ARM template.
In which order should you deploy the resources? To answer, move all resources from the list of resources to the answer area and arrange them in the correct order.

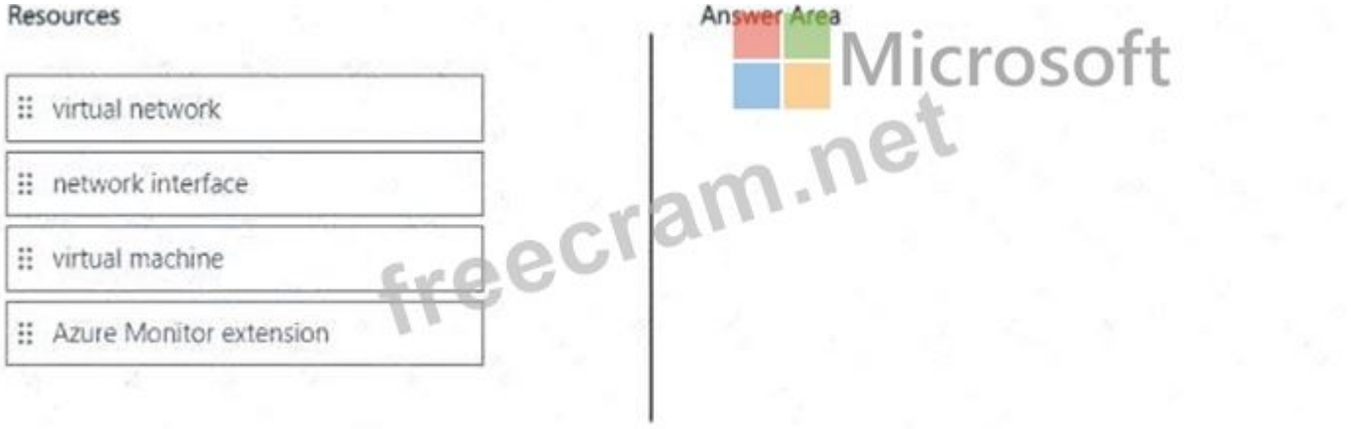

Answer:

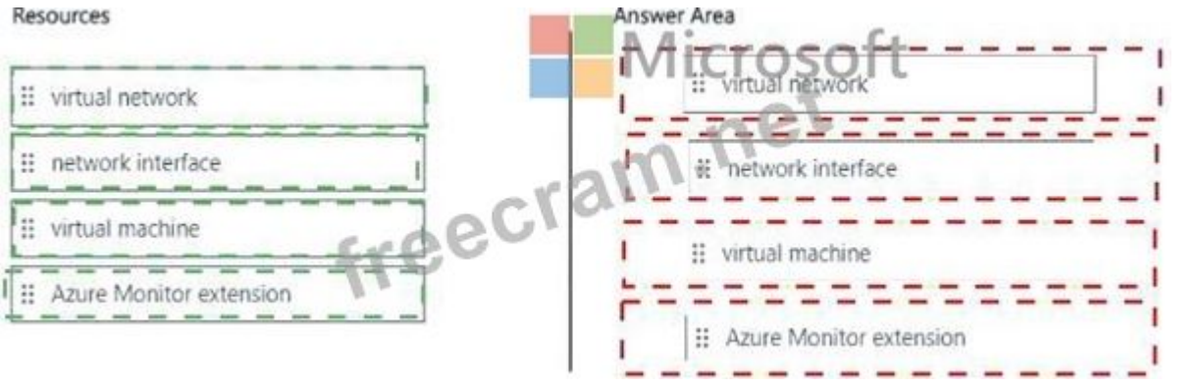

Explanation:

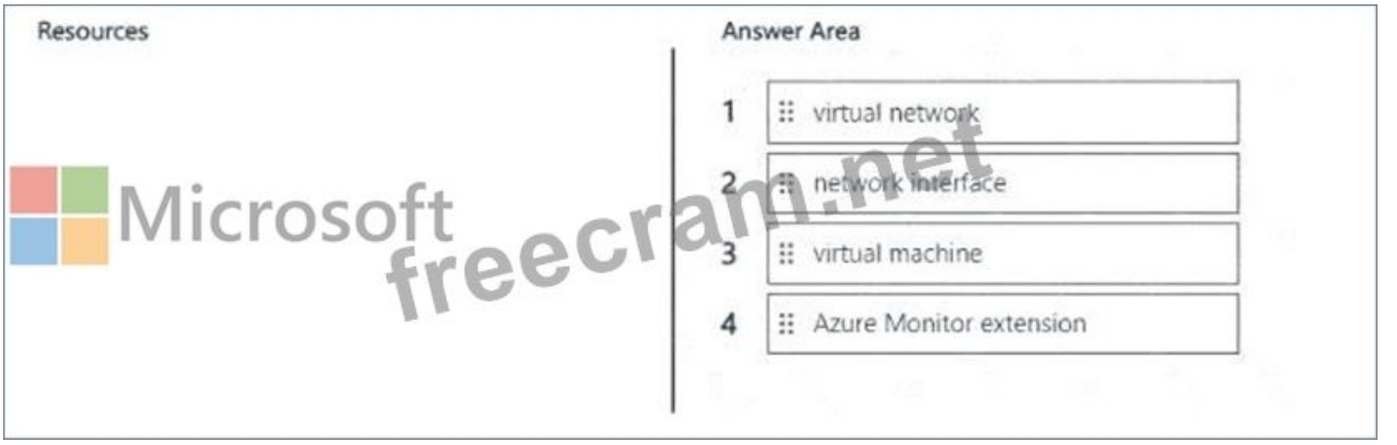

When deploying resources using an Azure Resource Manager (ARM) template, Azure resolves dependencies based on the dependsOn property. Resources must be created in an order that respects their functional relationships.

* Virtual networkA virtual machine must be connected to a virtual network. Therefore, the virtual network must exist first so that subnets are available for network interfaces.
* Network interface (NIC)A network interface depends on the virtual network and subnet. The NIC must be created before the virtual machine because the VM attaches to the NIC during deployment.
* Virtual machineThe virtual machine depends on the network interface. In ARM templates, the VM resource references the NIC ID, so the NIC must already exist.
* Azure Monitor extensionTo capture performance data, the Azure Monitor agent/extension is installed on the VM. Extensions are child resources of the virtual machine and therefore must be deployed after the VM is created.

Microsoft Azure Administrator documentation states that:

* Network interfaces require an existing virtual network.
* Virtual machines require an existing network interface.
* VM extensions depend on the virtual machine resource.

Final Verified Order:

Virtual network # Network interface # Virtual machine # Azure Monitor extension

NEW QUESTION: 83

You have an Azure subscription that contains three virtual networks named VNET1, VNET2, and VNET3.

Peering for VNET1 is configured as shown in the following exhibit.

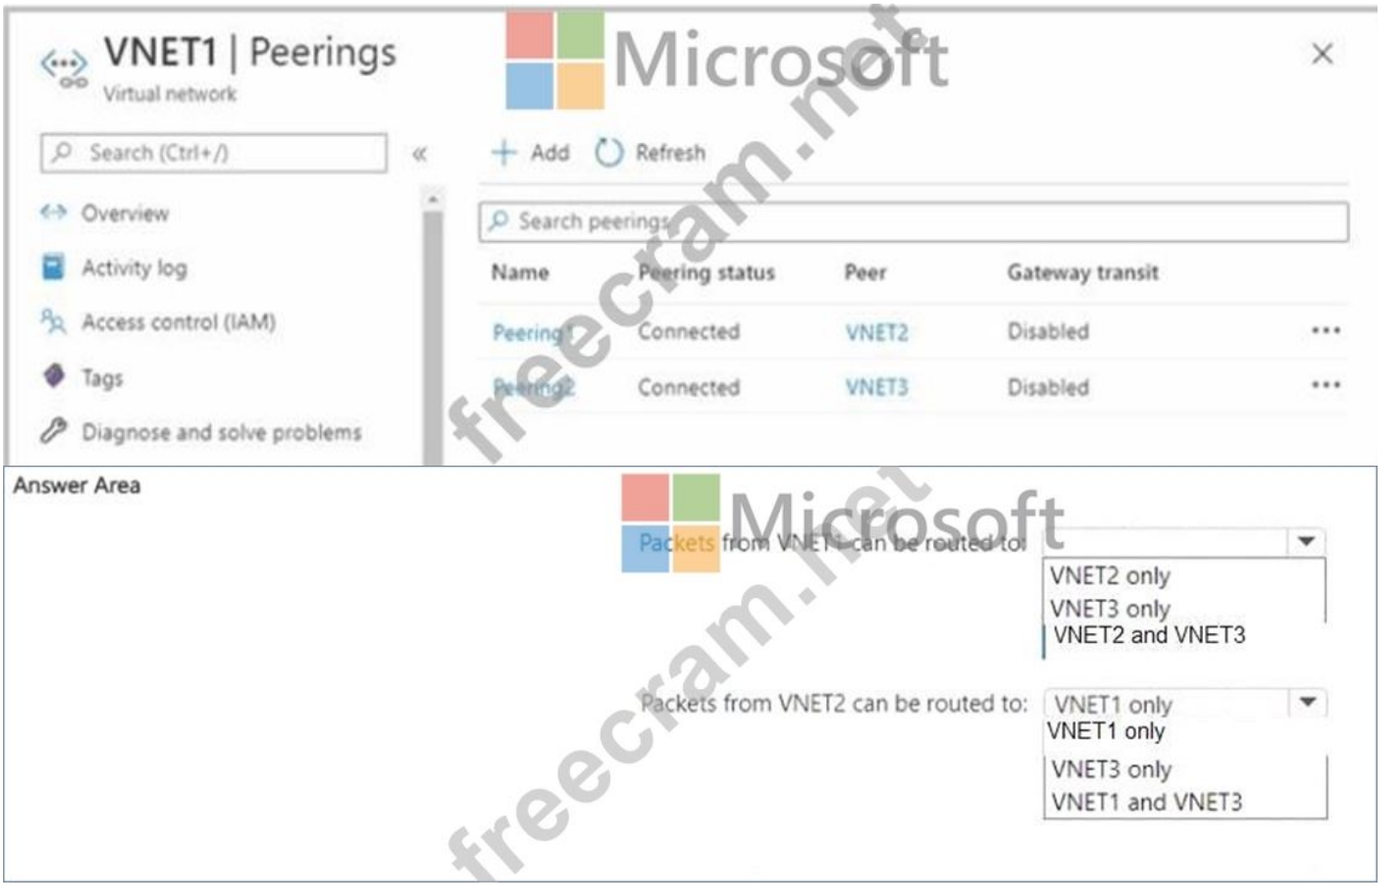

Answer:

Explanation:

This question evaluates understanding of Azure Virtual Network (VNet) peering behavior, specifically routing rules, non-transitive connectivity, and gateway transit settings.

From the exhibit, VNET1 has two peerings configured:

Peering1: VNET1 # VNET2 (Connected)

Peering2: VNET1 # VNET3 (Connected)

For both peerings, Gateway transit is Disabled.

Packets from VNET1

According to Microsoft Azure networking documentation:
VNet peering allows direct routing between peered VNets
A VNet can communicate with all VNets it is directly peered with
Since VNET1 is directly peered with both VNET2 and VNET3, traffic originating in VNET1 can be routed to both networks.
Packets from VNET1 can be routed to VNET2 and VNET3
Packets from VNET2
Azure VNet peering is not transitive. This means:
Even though VNET2 is peered with VNET1
And VNET1 is peered with VNET3
VNET2 does NOT automatically gain access to VNET3
Additionally:
Gateway transit is disabled
No user-defined routing or hub-and-spoke gateway configuration is present Therefore, VNET2 can only route traffic to its directly peered network, which is VNET1.
Packets from VNET2 can be routed to VNET1 only
Key Microsoft Azure Principles Applied
VNet peering enables direct connectivity only
Peering connections are non-transitive
Gateway transit must be explicitly enabled for transitive routing scenarios Without gateway transit or NVA routing, VNets cannot route through each other Final Verified Answer:
Packets from VNET1: # VNET2 and VNET3
Packets from VNET2: # VNET1 only

NEW QUESTION: 84
You implement the planned changes for Scope1.
You need to ensure that Scope1 meets the technical requirements.
What can you encrypt by using Scope1?
A. containers and blobs in storage2 only
B. containers and blobs in storage1 and storage2
C. containers, blobs, and file shares in storage2 only
D. containers, blobs, and file shares in storage1 and storage2
E. containers, blobs, file shares, queues, and tables in storage2 only
Answer: (SHOW ANSWER)
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: 85
You have an Anne container registry named Registry1 that contains an image named image1.
You receive an error message when you attempt to deploy a container instance by using image1.
You need to be able to deploy a container instance by using image1.
Solution: You assign the AcrPull role to ACR-Tasks-Network for Registry1.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
Azure Container Registry (ACR) uses Azure roles-based access control (RBAC) to determine who can pull, push, or manage images. The AcrPull role allows identities to pull (read) container images from a registry but not to push (write) or deploy images.
In this scenario, the error occurs when trying to deploy an Azure Container Instance (ACI) using an image from Registry1. The deployment process requires that the Azure Container Instance resource provider (the identity performing the deployment) has access permissions to pull the image from ACR. However, simply assigning the AcrPull role to an internal system identity like ACR-Tasks-Network does not necessarily authorize the ACI to pull images unless that specific managed identity used by ACI has the role assignment.
According to the Microsoft Azure Administrator documentation, the correct method to allow a container instance to pull an image from an Azure Container Registry is:
Enable a managed identity (system-assigned or user-assigned) for the Azure Container Instance.
Assign that managed identity the AcrPull role on the container registry.
The ACR Tasks Network identity mentioned in the scenario is used internally by ACR for automated build tasks, not for deployment of container instances. Therefore, assigning the AcrPull role to ACR-Tasks- Network does not resolve the issue of ACI being unable to access the container image.

Hence, this solution does not meet the goal because the role was assigned to the wrong identity. The correct approach is to assign AcrPull to the Azure Container Instance's managed identity instead.

NEW QUESTION: 86

You have an Azure subscription that contains the storage accounts shown in the following exhibit.

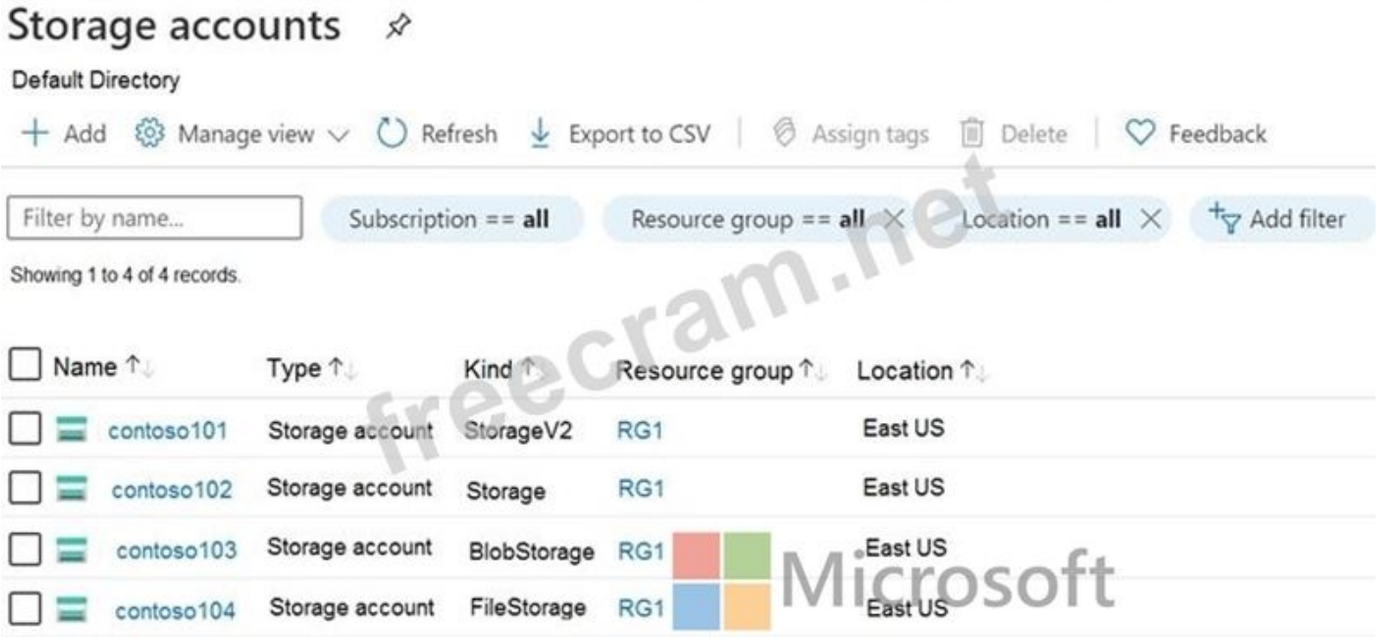

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

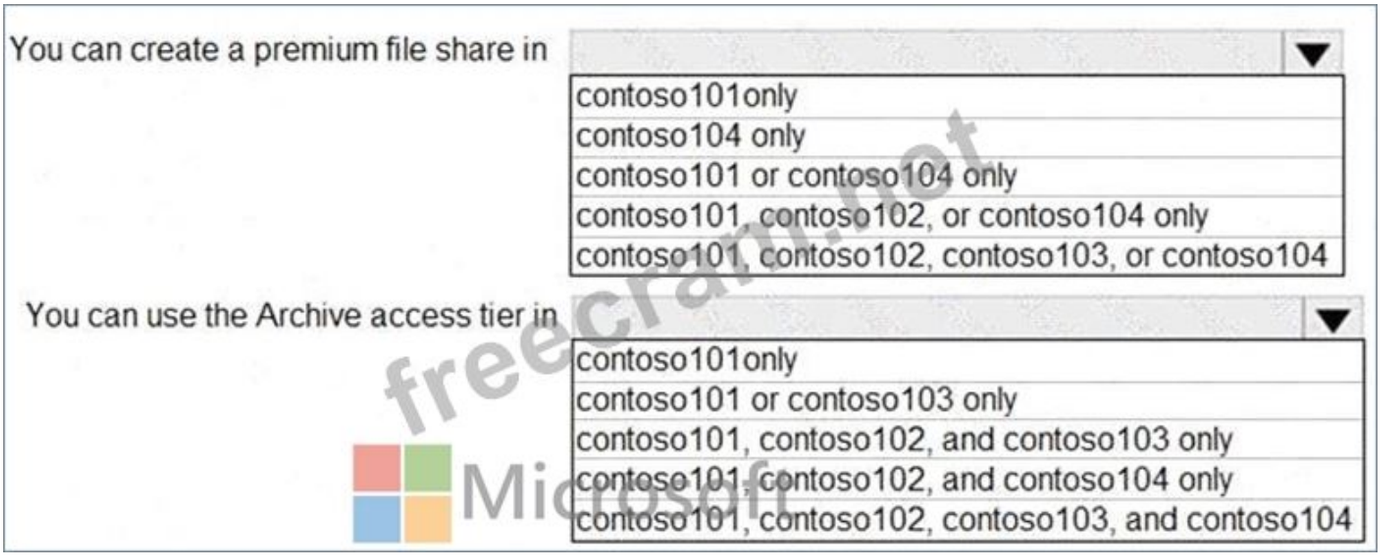

Answer:

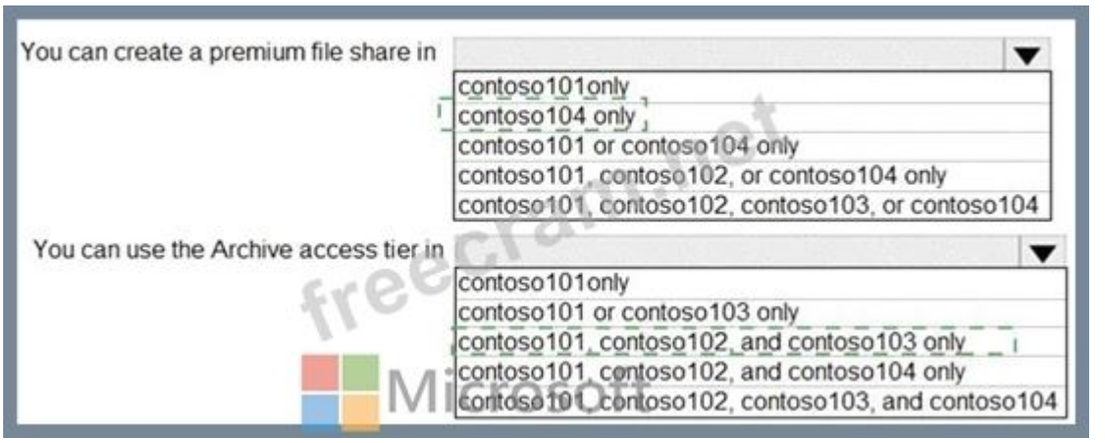

Explanation:

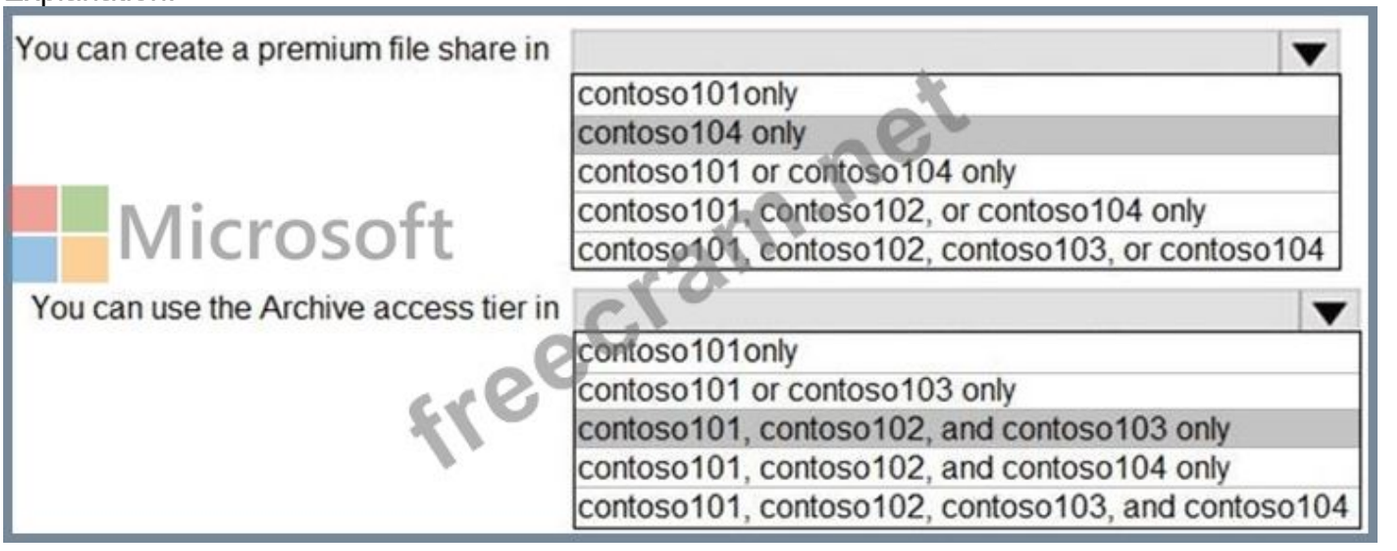

In Microsoft Azure Storage, the type and kind of a storage account determine which features and performance tiers are supported.

1## Premium File Share Support:

Premium file shares are only supported in FileStorage accounts, which provide high-performance, low-latency SSD-based storage optimized for enterprise file workloads. According to the official Azure Storage documentation (Microsoft Learn: "Azure Files performance tiers"), only the FileStorage account kind supports Premium file shares.

In the exhibit, only contoso104 is a FileStorage account. Therefore, premium file shares can only be created in contoso104.

2## Archive Access Tier Support:

The Archive access tier is available for Blob storage and General-purpose v2 (StorageV2) accounts that store block blobs. The Archive tier is used for data that is rarely accessed and provides the lowest storage cost.

From the exhibit:

* contoso101 (StorageV2) # supports archive tier.
* contoso102 (Storage) # older type, does not support access tiers beyond hot/cool.
* contoso103 (BlobStorage) # supports archive tier.
* contoso104 (FileStorage) # does not support Blob tiers.

Therefore, only contoso101, contoso102, and contoso103 allow the Archive tier, but since

standard "Storage" (contoso102) doesn't support Archive, the correct accounts are contoso101 and contoso103.

NEW QUESTION: 87

You have an Azure subscription that contains the virtual machines shown in the following table.

javascript:void(0)

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

You deploy a load balancer that has the following configurations:

* Name: LB1
* Type internal
* SKU: Standard
* Virtual network VNET1

You need to ensure that you can add VM1 and VM2 to the backend pool of LB1.

Solution: You create a Basic SKU public IP address, associate the address to the network interface of VM1, and then start VM1.

Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

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: 88

You have two Azure subscriptions named Sub1 and Sub2.

Sub1 contains a virtual machine named VM1 and a storage account named storage1.

VM1 is associated to the resources shown in the following table.

You need to move VM1 to Sub2.

Which resources should you move to Sub2?

A. VM1, Disk1. and NetInt1 only

B. VM1. Disk1. and VNet1 only

C. VM1. Disk1. and storage1 only

D. VM1. Disk1. NetInt1, and VNet1

Answer: (SHOW ANSWER)

When you move a virtual machine to a different subscription, you need to move all the resources that are associated with the virtual machine, such as the disks, the network interface, and the virtual network. You cannot move a virtual machine without moving its dependent resources. You also need to ensure that the target subscription supports the same region, resource type, and API version as the source subscription. Then, References: [Move a Windows VM to another Azure subscription or resource group]

NEW QUESTION: 89

You have an Azure App Service app named WebApp1 that contains two folders named Folder1 and Folder2.

You need to configure a daily backup of WebApp1. The solution must ensure that Folder2 is excluded from the backup.

What should you create first and what should you use to exclude Fokier2? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

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: 90
You have an Azure subscription that contains the Azure virtual machines shown in the following table.

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

You configure the network interfaces of the virtual machines to use the settings shown in the following table

Name	DNS server
VM1	*None*
VM2	192.168.10.15
VM3	192.168.10.15

From the settings of VNET1, you configure the DNS servers shown in the following exhibit.

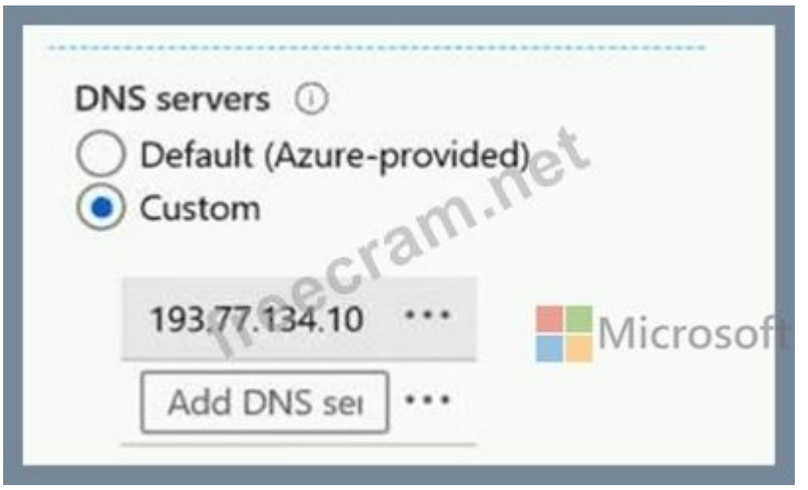

The virtual machines can successfully connect to the DNS server that has an IP address of 192.168.10.15 and the DNS server that has an IP address of 193.77.134.10.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.

	Yes	No
VM1 connects to 193.77.134.10 for DNS queries.	○	○
VM2 connects to 193.77.134.10 for DNS queries.	○	○
VM3 connects to 192.168.10.15 for DNS queries.	○	○

Answer:

Explanation:

Statements	Yes	No
VM1 connects to 193.77.134.10 for DNS queries.	○	○
VM2 connects to 193.77.134.10 for DNS queries.	○	○
VM3 connects to 192.168.10.15 for DNS queries.	○	○

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: 91
You have an Azure subscription that contains a web app named webapp1. You need to add a custom domain named www.contoso.com to webapp1. What should you do first?
A. Upload a certificate.
B. Add a connection string.
C. Stop webapp1.
D. Create a DNS record.
Answer: (SHOW ANSWER)
When configuring a custom domain for an Azure Web App (App Service) such as webapp1, the first prerequisite is to create and configure a corresponding DNS record that maps your custom domain name (www.contoso.com) to the Azure App Service's default domain (webapp1.azurewebsites.net).
According to Microsoft Azure App Service documentation:
"Before you can assign a custom domain name to your web app, you must first create a CNAME record in your DNS provider that maps your custom domain (for example, www.contoso.com) to your app's default domain name (webappname.azurewebsites.net)." This DNS record ensures that requests to the custom domain are routed correctly to Azure's front-end load balancer

hosting your web app. Only after the DNS record has propagated can you validate and bind the domain to the web app in the Azure portal.
Uploading an SSL certificate is only required for HTTPS bindings after the domain is added.
Stopping the app or adding connection strings is unrelated to domain configuration.
Therefore, the first step is to create a DNS record at your domain registrar that points to your web app.

NEW QUESTION: 92
You develop the following Azure Resource Manager (ARM) template to create a resource group and deploy an Azure Storage account to the resource group.
Which cmdlet should you run to deploy the template?
A. New-AzTenantDeployment
B. New-AzResourceGroupDeployment
C. New-AzResource
D. New-AzOeployment
Answer: (SHOW ANSWER)
In Microsoft Azure, Azure Resource Manager (ARM) templates are JSON-based files used to define and automate the deployment of Azure resources in a consistent and repeatable manner. To deploy an ARM template, the PowerShell cmdlet you use depends on the deployment scope - whether the template targets a tenant, management group, subscription, or resource group.
According to the Microsoft Azure Administrator documentation (AZ-104 Study Guide, "Deploy and Manage Azure Resource Manager Templates"):
New-AzTenantDeployment # Used for tenant-level deployments, typically involving management group creation or tenant-wide policies.
New-AzSubscriptionDeployment # Used for subscription-level deployments, when resources are created directly under the subscription rather than within a resource group.
New-AzResourceGroupDeployment # Used for resource group-level deployments, which is the most common scenario for creating and managing resources such as virtual machines, storage accounts, and virtual networks.
New-AzDeployment # An older alias that may refer to subscription-level deployments in older Az modules but is not the standard for resource group deployments.
In this question, the ARM template defines both a resource group and a storage account that will be deployed within that resource group. Therefore, the correct PowerShell cmdlet to use is:
New-AzResourceGroupDeployment `

-ResourceGroupName <RGName> `
-TemplateFile <PathToTemplateFile>
This cmdlet processes the specified ARM template file, validates the schema, and provisions all the resources defined in the template inside the target resource group.
Hence, based on the Azure Resource Manager official documentation and exam objectives, the correct answer is:

NEW QUESTION: 93
You have an Azure subscription that contains the storage accounts shown in the following table.

Name	Kind	Performance	Replication	Access tier
Storage1	Storage (general purpose v1)	Premium	Geo-redundant storage (GRS)	*None*
Storage2	StorageV2 (general purpose v2)	Standard	Locally-redundant storage (LRS)	Cool
Storage3	StorageV2 (general purpose v2)	Premium	Read-access geo-redundant storage (RA-GRS)	Hot
Storage4	BlobStorage	Standard	Locally-redundant storage (LRS)	Hot

You need to identify which storage account can be converted to zone-redundant storage (ZRS) replication by requesting a live migration from Azure support.
What should you identify?
A. Storage1
B. Storage2
C. Storage3
D. Storage4
Answer: (SHOW ANSWER)
https://learn.microsoft.com/en-us/azure/storage/common/redundancy-migration?tabs=portal

NEW QUESTION: 94
You have two Azure virtual machines named VM1 and VM2 that run Windows Server. The virtual machines are in a subnet named Subnet1. Subnet1 is in a virtual network named VNet1. You need to prevent VM1 from accessing VM2 on port 3389.
What should you do?
A. Create a network security group (NSG) that has an outbound security rule to deny destination port 3389 and apply the NSG to the network interface of VM1.
B. Create a network security group (NSG) that has an inbound security rule to deny source port 3389 and apply the NSG to Subnet1.
C. Configure Azure Bastion in VNet1.
D. Create a network security group (NSG) that has an outbound security rule to deny source port 3389 and apply the NSG to Subnet1.
Answer: (SHOW ANSWER)

NEW QUESTION: 95

You have an Azure subscription named Subscription1 that contains the storage accounts shown

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

You plan to use the Azure Import/Export service to export data from Subscription1.

Which account can be used to export the data.

What should you identify?

A. storage1

B. storage2

C. storage3

D. storage4

Answer: (SHOW ANSWER)

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: 96

You have an Azure subscription that contains an Azure Availability Set named WEBPROD-AS-USE2 as shown in the following exhibit.

You add 14 virtual machines to WEBPROD-AS-USE2.

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

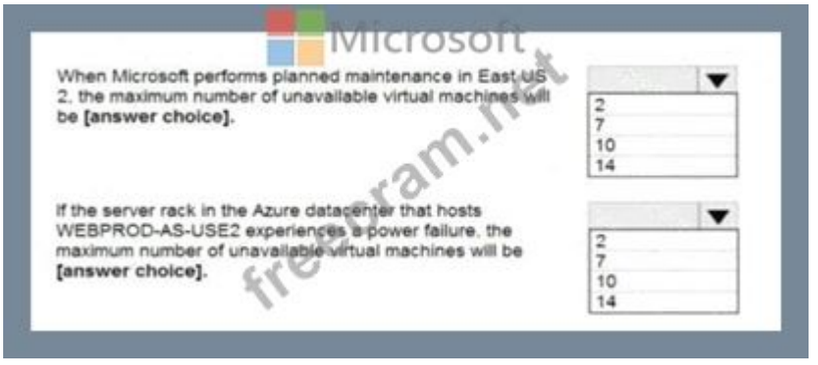

Answer:

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.

Answer Area

When Microsoft performs planned maintenance in East US 2, the maximum number of unavailable virtual machines will be **[answer choice].**

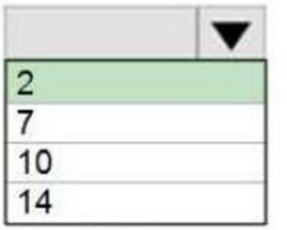

If the server rack in the Azure datacenter that hosts WEBPROD-AS-USE2 experiences a power failure, the maximum number of unavailable virtual machines will be **[answer choice].**

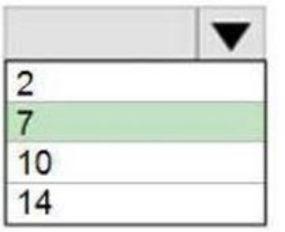

Reference:
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/manage-availability

NEW QUESTION: 97
You have an Azure subscription that contains 10 virtual networks. The virtual networks are hosted in separate resource groups.
Another administrator plans to create several network security groups (NSGs) in the subscription.
You need to ensure that when an NSG is created, it automatically blocks TCP port 8080 between the virtual networks.
Solution: You configure a custom policy definition, and then you assign the Azure policy to the subscription.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In Azure, Azure Policy is a governance tool used to enforce organizational standards and assess compliance across Azure resources. It allows administrators to define and assign policy definitions that automatically audit, deny, or modify resource configurations at deployment or runtime.
In this scenario, the requirement is that every time a Network Security Group (NSG) is created, it should automatically contain a rule that blocks TCP port 8080 traffic between virtual networks.

The Microsoft Azure Policy documentation confirms that you can create a custom policy definition using the Microsoft.Network/networkSecurityGroups/securityRules resource type. Within the policy's JSON definition, you can specify conditions such as:
* The resource type to which the policy applies (networkSecurityGroups).
* The enforcement mode (deny or deployIfNotExists).
* The required configuration, such as a specific inbound or outbound rule (in this case, a rule denying TCP 8080).
By using the DeployIfNotExists effect in the policy, Azure automatically ensures that the NSG includes the required rule. If the rule does not exist, Azure deploys it automatically during resource creation.
Assigning this custom policy definition at the subscription level ensures it is inherited by all resource groups and applies to all virtual networks created in that subscription. This meets the goal because the requirement is to enforce a security configuration across all NSGs, regardless of which resource group or virtual network they belong to.
Therefore, configuring and assigning a custom Azure Policy to the subscription fully satisfies the requirement.

NEW QUESTION: 98
You deploy an Azure Kubernetes Service (AKS) cluster that has the network profile shown in the following exhibit.

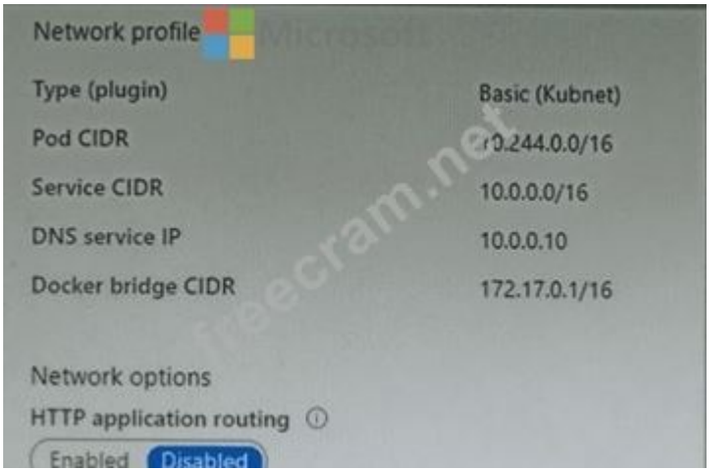

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

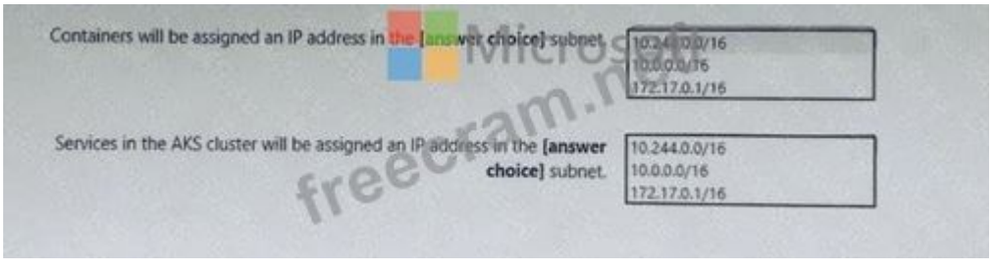

Answer:

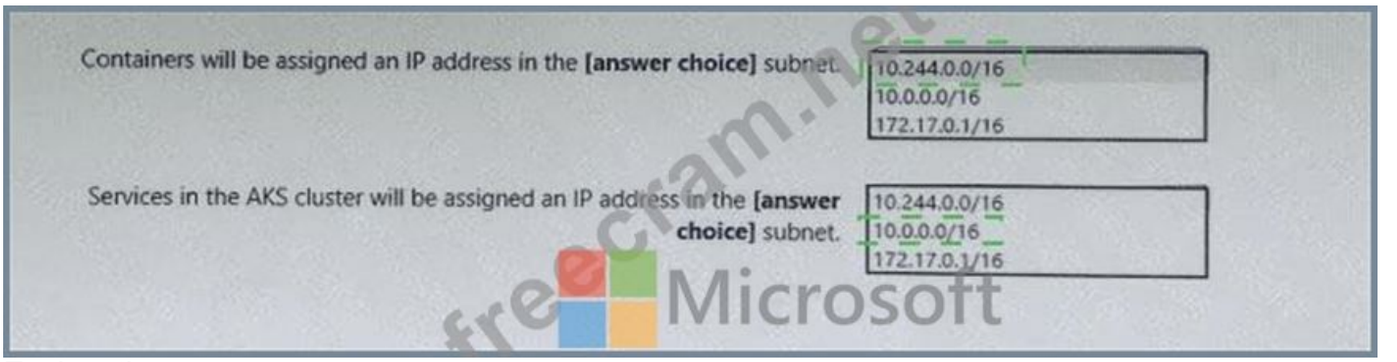

Explanation:

Azure Kubernetes Service (AKS) allows administrators to choose between two primary network models for pods and services - Kubenet (Basic) and Azure CNI (Advanced).

In the Basic (Kubenet) networking model - as shown in the image - AKS manages a separate internal network space for pods and services distinct from the virtual network (VNet) of the nodes.

Here's how each CIDR range functions in the network profile:

Pod CIDR (10.244.0.0/16)

This range is used by Kubernetes to assign IP addresses to individual containers (pods) within the cluster.

Each node gets a subset of IPs from this range, and pods on that node use these IPs for internal cluster communication.

These pod IPs are not routable outside the cluster unless Network Address Translation (NAT) is applied.

Service CIDR (10.0.0.0/16)

This subnet is used by Kubernetes to allocate virtual IP addresses for services (e.g., ClusterIP).

These IPs represent internal load balancers that direct traffic to the backend pods through kube-proxy.

The DNS Service IP (10.0.0.10) resides in this range, representing the internal DNS service (kube-dns or CoreDNS).

Docker Bridge CIDR (172.17.0.1/16)

This range is used for the Docker bridge network within each node to manage local container networking before being connected to Kubernetes.

It is not used for pods or services within AKS.

Therefore:

Containers (Pods) # IPs assigned from 10.244.0.0/16 (Pod CIDR).

Services # IPs assigned from 10.0.0.0/16 (Service CIDR).

These CIDRs are critical when integrating AKS clusters with other Azure VNets to avoid overlapping address spaces and ensure seamless connectivity.

NEW QUESTION: 99

You have two Azure subscriptions named Sub1 and Sub2. Sub1 is in a management group named MG1. Sub2 is in a management group named MG2.

You have the resource groups shown in the following table.

Name	Subscription
RG1	Sub1
RG2	Sub2

You have the virtual machines shown in the following table.

Name	Resource group
VM1	RG1
VM2	RG2
VM3	RG2

You assign roles to users as shown in the following table.

User	Role	Resource
User1	Virtual Machine Contributor	MG1
User1	Virtual Machine User Login	Sub2
User2	Virtual Machine Contributor	MG2
User2	Virtual Machine User Login	Sub1
User2	Virtual Machine User Login	VM3

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
User1 can sign in to VM1.	○	○
User2 can manage disks and disk snapshots of VM1.	○	○
User2 can manage disks and disk snapshots of VM3.	○	○

Answer:

Answer Area

Statements	Yes	No
User1 can sign in to VM1.	○	○
User2 can manage disks and disk snapshots of VM1.	○	○
User2 can manage disks and disk snapshots of VM3.	○	○

Explanation:
User 1 can sign in to VMI. = YES
User 1 has the Virtual Machine User Login role assigned at the scope of RG1. This role allows the user to sign in to virtual machines in the resource group using Azure AD credentials. VMI is a virtual machine in RG1, so User 1 can sign in to it.
User 2 can manage disks and disk snapshots of VMI. = NO
User 2 has the Disk Snapshot Contributor role assigned at the scope of MG2. This role allows the user to manage disk snapshots in the management group. However, VMI is not in MG2, but in RG1, which is in MG1. Therefore, User 2 does not have the permission to manage disks and disk snapshots of VMI.
User 2 can manage disks and disk snapshots of VM3. = YES
User 2 has the Disk Snapshot Contributor role assigned at the scope of MG2. This role allows the user to manage disk snapshots in the management group. VM3 is a virtual machine in RG3, which is in Sub2, which is in MG2. Therefore, User 2 has the permission to manage disks and disk snapshots of VM3.

NEW QUESTION: 100
You have an Azure subscription.
You plan to deploy a storage account named storage' by using the following Azure Resource Manager (ARM) template.

```
{
    "$schema": "http://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
    "contentVersion": "1.0.0.0",
    "resources": [
        {
            "name": "storage1",
            "type": "Microsoft.Storage/storageAccounts",
            "apiVersion": "2021-08-01",
            "location": "East US",
            "properties": {
                "allowBlobPublicAccess": true,
                "defaultToOAuthAuthentication": false,
                "networkAcls": {
                    "bypass": "AzureServices",
                    "defaultAction": "Allow",
                    "ipRules": []
                },
                "isVersioningEnabled": true
            },
            "dependsOn": [
                "[concat('Microsoft.Storage/storageAccounts/', 'storage1')]"
            ]
        }
    ]
}
```


For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer:

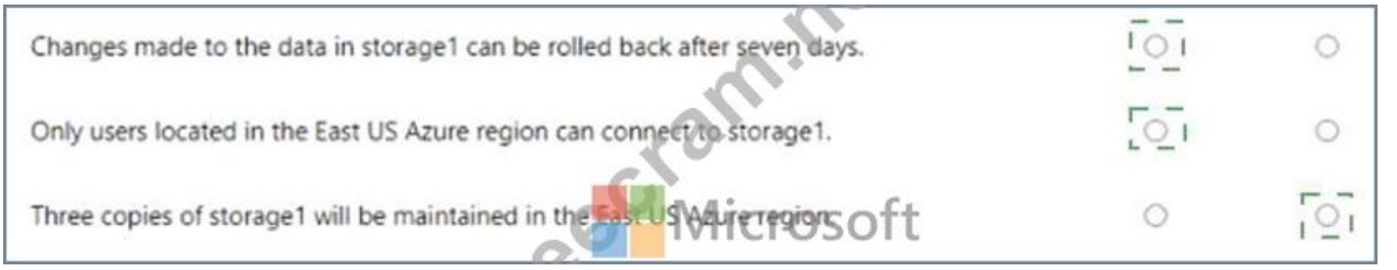

Explanation:

NO

NO

YES

The provided ARM template defines an Azure Storage Account resource with the following key properties:

```
{
"type": "Microsoft.Storage/storageAccounts",
"name": "storage1",
"location": "East US",
"properties": {
"allowBlobPublicAccess": true,
"defaultToOAuthAuthentication": false,
"networkAcls": {
"bypass": "AzureServices",
"defaultAction": "Allow",
"ipRules": []
},
"isVersioningEnabled": true
}
}
```

Let's analyze each statement based on Azure Storage behavior and documentation:

Statement 1: "Changes made to the data in storage1 can be rolled back after seven days."

Incorrect

The property "isVersioningEnabled": true refers to blob versioning, which maintains previous versions of objects (blobs) when they are modified or deleted. However, there is no automatic seven-day rollback period.

Versions remain indefinitely until manually deleted or managed via lifecycle policies.

Microsoft documentation:

"When blob versioning is enabled, Azure Storage automatically maintains previous versions of an object. You can restore a previous version manually, but there is no automatic rollback period." Hence, changes can be rolled back manually, not automatically after seven days.
Statement 2: "Only users located in the East US Azure region can connect to storage1."
Incorrect
The networkAcls section specifies:
"defaultAction": "Allow",
"bypass": "AzureServices",
"ipRules": []
This configuration means all network access is allowed by default ("defaultAction": "Allow"). There are no IP restrictions, so users from any region can connect.
Azure documentation states:
"If the defaultAction is set to Allow, traffic from all networks can access the storage account unless specific network rules are added." Therefore, access is not restricted to the East US region.
Statement 3: "Three copies of storage1 will be maintained in the East US Azure region."
Correct
By default, if redundancy options like ZRS or GRS are not specified, Azure Storage uses Locally Redundant Storage (LRS).
From the Microsoft study guide:
"Locally Redundant Storage (LRS) maintains three synchronous copies of your data within a single datacenter in the primary region." Since the location is "East US" and no redundancy property is explicitly defined, the default LRS applies - meaning three copies within the same region (East US).

NEW QUESTION: 101
You are planning the move of App1 to Azure.
You create a network security group (NSG).
You need to recommend a solution to provide users with access to App1.
What should you recommend?
A. Create an outgoing security rule for port 443 from the Internet. Associate the NSG to all the subnets.
B. Create an incoming security rule for port 443 from the Internet. Associate the NSG to all the subnets.
C. Create an incoming security rule for port 443 from the Internet. Associate the NSG to the subnet that contains the web servers.
D. Create an outgoing security rule for port 443 from the Internet. Associate the NSG to the subnet that contains the web servers.
Answer: (SHOW ANSWER)
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: 102
You have an Azure subscription that contains three virtual machines named VM1, VM2. and VM3. All the virtual machines are in an availability set named AVSetl. You need to scale up VM1 to a new virtual machine size, but the intended size is unavailable. What should you do first?
A. Shut down VM2 and VM3.
B. Convert AVSetl into a managed availability set.
C. DeallocateVM1.
D. Create a proximity placement group.
Answer: (SHOW ANSWER)
When resizing a virtual machine in Azure, the target size must be available on the same hardware cluster where the VM currently resides. If the desired size is unavailable, Azure prevents the resize operation.
According to the Microsoft Azure Administrator guide:
"To resize a virtual machine to a size not available on the current hardware cluster, the VM must first be deallocated. When a VM is deallocated, its current hardware association is released, allowing Azure to place it on a different cluster that supports the new size." In this scenario, VM1 is part of an availability set (AVSet1). All VMs in an availability set must reside on the same hardware cluster. Therefore, to resize VM1 successfully to a new size that isn't currently available, you must first deallocate the VM, which frees it from its existing cluster allocation. After deallocation, you can resize it, and Azure will allocate it to a compatible cluster.
NEW QUESTION: 103
You have an Azure Active Directory (Azure AD) tenant named contoso.com.
You have a CSV file that contains the names and email addresses of 500 external users.

You need to create a guest user account in contoso.com for each of the 500 external users.
Solution: You create a Power Shell script that runs the New-MgUser cmdlet for each user.
Does this meet the goal?

A. Yes

B. NO

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: 104

You plan to deploy five virtual machines to a virtual network subnet.
Each virtual machine will have a public IP address and a private IP address.
Each virtual machine requires the same inbound and outbound security rules.
What is the minimum number of network interfaces and network security groups that you require?
To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

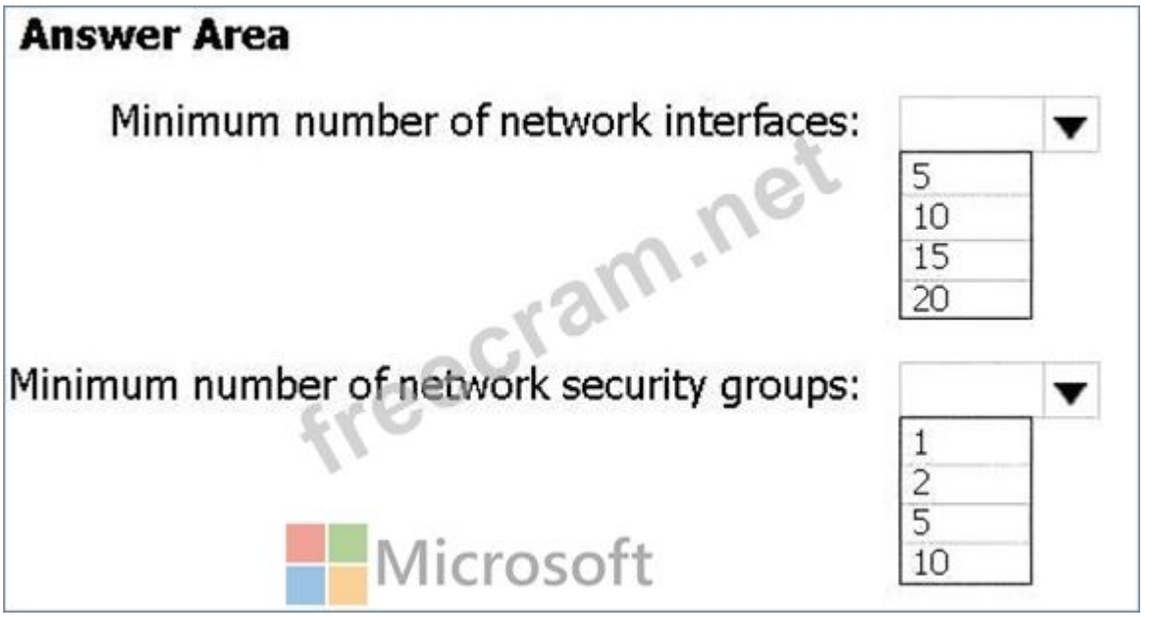

Answer:

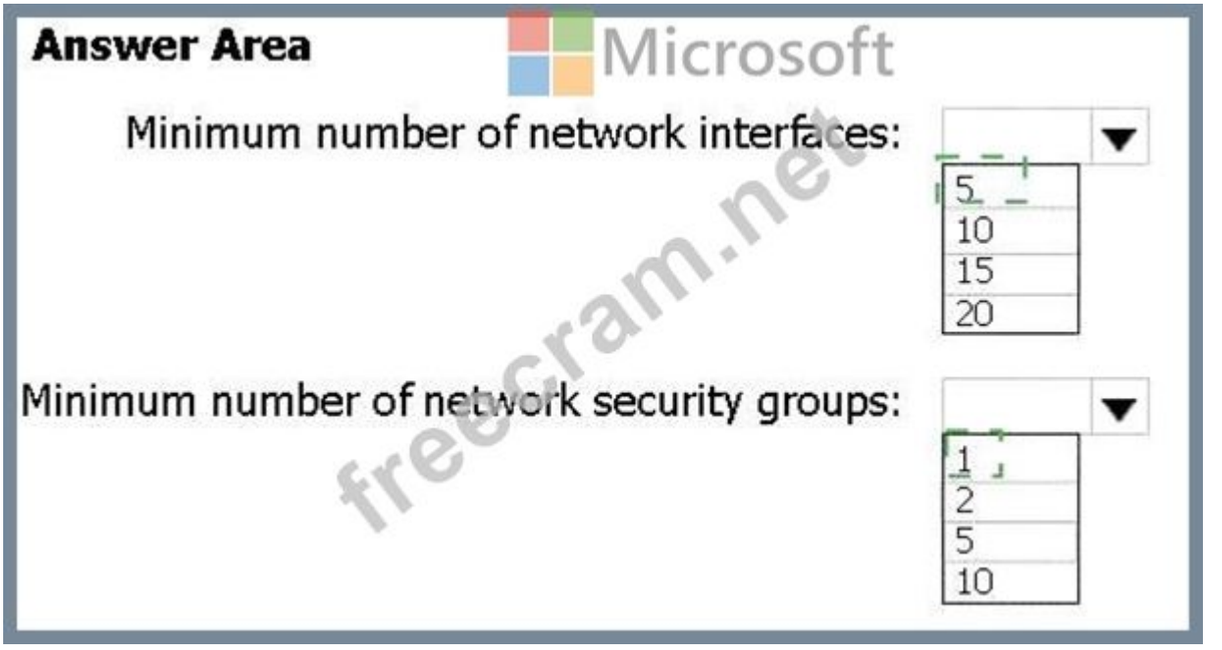

Explanation:

Box 1: 5

A public and a private IP address can be assigned to a single network interface.

Box 2: 1

You can associate zero, or one, network security group to each virtual network subnet and network interface in a virtual machine. The same network security group can be associated to as many subnets and network interfaces as you choose.

Answer Area

Reference:
https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface-addresses

NEW QUESTION: 105
You have an Azure subscription that contains The storage accounts shown in the following table.

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

You deploy a web app named Appl to the West US Azure region.
You need to back up Appl. The solution must minimize costs.
Which storage account should you use as the target for the backup?
A. storage1
B. storage2
C. storage3
D. storage4
Answer: (SHOW ANSWER)
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: 106

You have an Azure subscription named Subscription! that contains the resources shown in the following table.

Name	Type	Region	Resource group
RG1	Resource group	West US	*Not applicable*
RG2	Resource group	West US	*Not applicable*
Vault1	Recovery Services vault	Central US	RG1
Vault2	Recovery Services vault	West US	RG2
VM1	Virtual machine	Central US	RG2
storage1	Storage account	West US	RG1
SQL1	Azure SQL database	East US	RG2

In storage1, you create a blob container named blob1 and a file share named share1.

Which resources can be backed up to Vault1 and Vault2? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

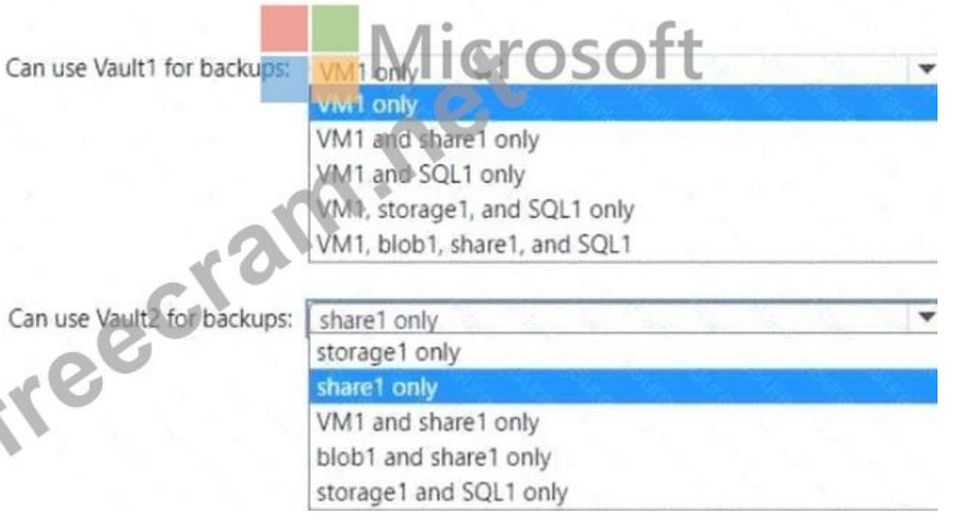

Answer:

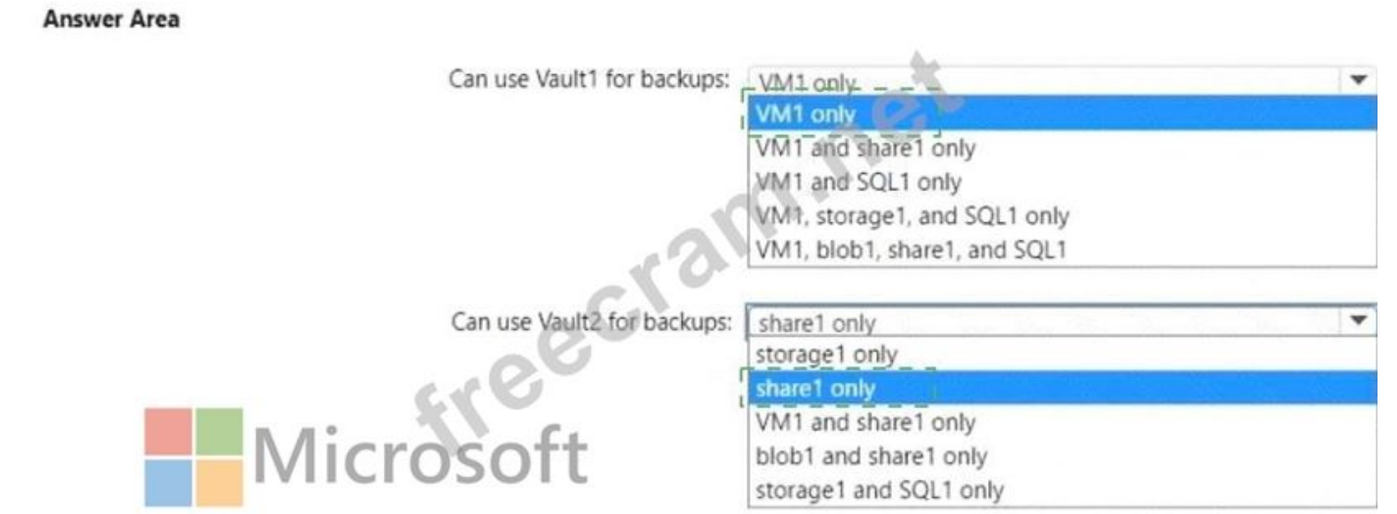

Explanation:

Azure Recovery Services vaults are used to store backup data and recovery points for supported workloads such as Azure VMs, Azure Files, and on-premises workloads. However, the configuration rules are strictly defined based on region and supported resource type.

From Microsoft Azure Administrator documentation and the official Azure Backup architecture, the following principles apply:

* Backup Region Dependency - A Recovery Services vault can only back up resources located in the same Azure region. For example, a vault in Central US can only back up Azure VMs, storage accounts, or workloads in Central US. It cannot back up resources in West US or East US.

* Vault1 (Central US, RG1)

* Located in Central US and can therefore back up only resources in Central US.

* From the table: VM1 is in Central US, and is eligible.

* storage1 (West US) and SQL1 (East US) cannot be backed up to Vault1 due to cross-region restrictions.

Therefore, Vault1 can back up VM1 only.

* Vault2 (West US, RG2)

* Located in West US, and can only back up resources in West US.

* In West US, the only relevant resource is storage1, which contains a file share named share1. Azure Backup supports Azure File share backup via Recovery Services vaults.

* Vault2 cannot back up VM1 (Central US) or SQL1 (East US) since they reside in different regions.

Therefore, Vault2 can back up share1 only.

* Azure Backup Limitations Summary (From Microsoft Docs):

* Azure VMs # must be backed up to a vault in the same region.

* Azure File Shares # can be backed up using Recovery Services vaults in the same region as the storage account.

* Azure SQL Database # uses its own automated backup mechanism, not Recovery Services vaults.
Hence, by Azure documentation compliance and AZ-104 study guide principles:
* Vault1 backs up VM1 only.
* Vault2 backs up share1 only.
Final Verified Answer:
Vault1 - VM1 only
Vault2 - share1 only

NEW QUESTION: 107
You have Azure subscriptions named Subscription1 and Subscription2.
Subscription1 has following resource groups:

Name	Region	Lock type
RG1	West Europe	None
RG2	West Europe	Read Only

RG1 includes a web app named App1 in the West Europe location.
Subscription2 contains the following resource groups:

Name	Region	Lock type
RG3	East Europe	Delete
RG4	Central US	none

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

App1 present in RG1 and in RG1 there is no lock available. So you can move App1 to other resource groups, RG2, RG3, RG4.

Note:

App Service resources can only be moved from the resource group in which they were originally created. If an App Service resource is no longer in its original resource group, move it back to its original resource group.

Reference:

https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/move-limitations/app-service- move-limitations

NEW QUESTION: 108
You have an Azure subscription that contains two storage accounts named contoso101 and contoso102.
The subscription contains the virtual machines shown in the following table.
VNet1 has service endpoints configured as shown in the Service endpoints exhibit. (Click the Service endpoints tab.)

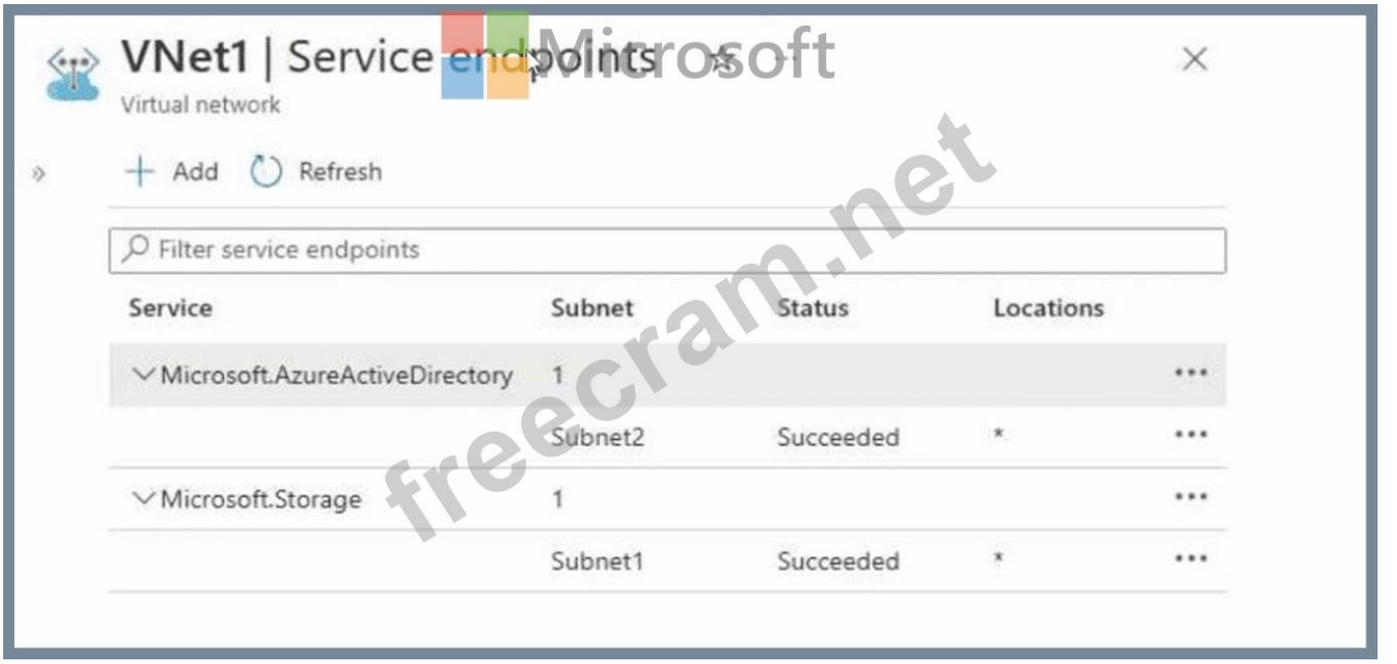

The Microsoft. Storage service endpoint has the service endpoint policy shown in the Microsoft. Storage exhibit. (Click the Microsoft. Storage tab.)

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:
Answer Area

Explanation:

NEW QUESTION: 109

You have an Azure subscription that contains a resource group named RG1.

You need to prevent administrators from inadvertently modifying the resources in RG1.

How should you complete the PowerShell command? To answer, select the options in the answer area.

NOTE: Each correct answer is worth one point.

Answer:

Explanation:

Azure provides resource locks to prevent accidental changes or deletion of critical resources. This is a core governance feature covered in the Microsoft Azure Administrator (AZ-104) study materials under resource management and protection.

To prevent administrators from inadvertently modifying resources in a resource group, you must apply a ReadOnly lock at the resource group scope. A ReadOnly lock allows users to view resources but blocks any modification or deletion, even for users with high privileges, unless the lock is first removed.

There are two primary lock levels in Azure:

ReadOnly - Users can read resources but cannot modify or delete them.

CanNotDelete (Delete) - Users can modify resources but cannot delete them.

Since the requirement is to prevent modification, the correct lock level is ReadOnly, not DeleteResources or CanNotDelete.

From Microsoft Azure Administrator documentation:

"A ReadOnly lock prevents users from making changes to a resource. Authorized users can still

remove the lock if needed." To create this lock using PowerShell, the correct cmdlet is New-AzResourceLock, which is used to create management locks at the subscription, resource group, or resource level.
Applying the lock to RG1 ensures all resources within the group inherit the protection, meeting the requirement with minimal administrative effort.
Correct PowerShell Command (Conceptual Form)
New-AzResourceLock `
-LockName LockGroup `
-LockLevel ReadOnly `
-ResourceGroupName RG1
Final Hotspot Selections
Cmdlet: New-AzResourceLock
LockName: LockGroup
LockLevel: ReadOnly
ResourceGroupName: RG1

NEW QUESTION: 110

You need to recommend an identify solution that meets the technical requirements.
What should you recommend?

A. federated single-on (SSO) and Active Directory Federation Services (AD FS)
B. password hash synchronization and single sign-on (SSO)
C. cloud-only user accounts
D. Pass-through Authentication and single sign-on (SSO)
Answer: (SHOW ANSWER)
Active Directory Federation Services is a feature and web service in the Windows Server Operating System that allows sharing of identity information outside a company's network.
Scenario: Technical Requirements include:
Prevent user passwords or hashes of passwords from being stored in Azure.
References: https://www.sherweb.com/blog/active-directory-federation-services/

NEW QUESTION: 111
You need to create storage5. The solution must support the planned changes.
Which type of storage account should you use, and which account should you configure as the destination storage account? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

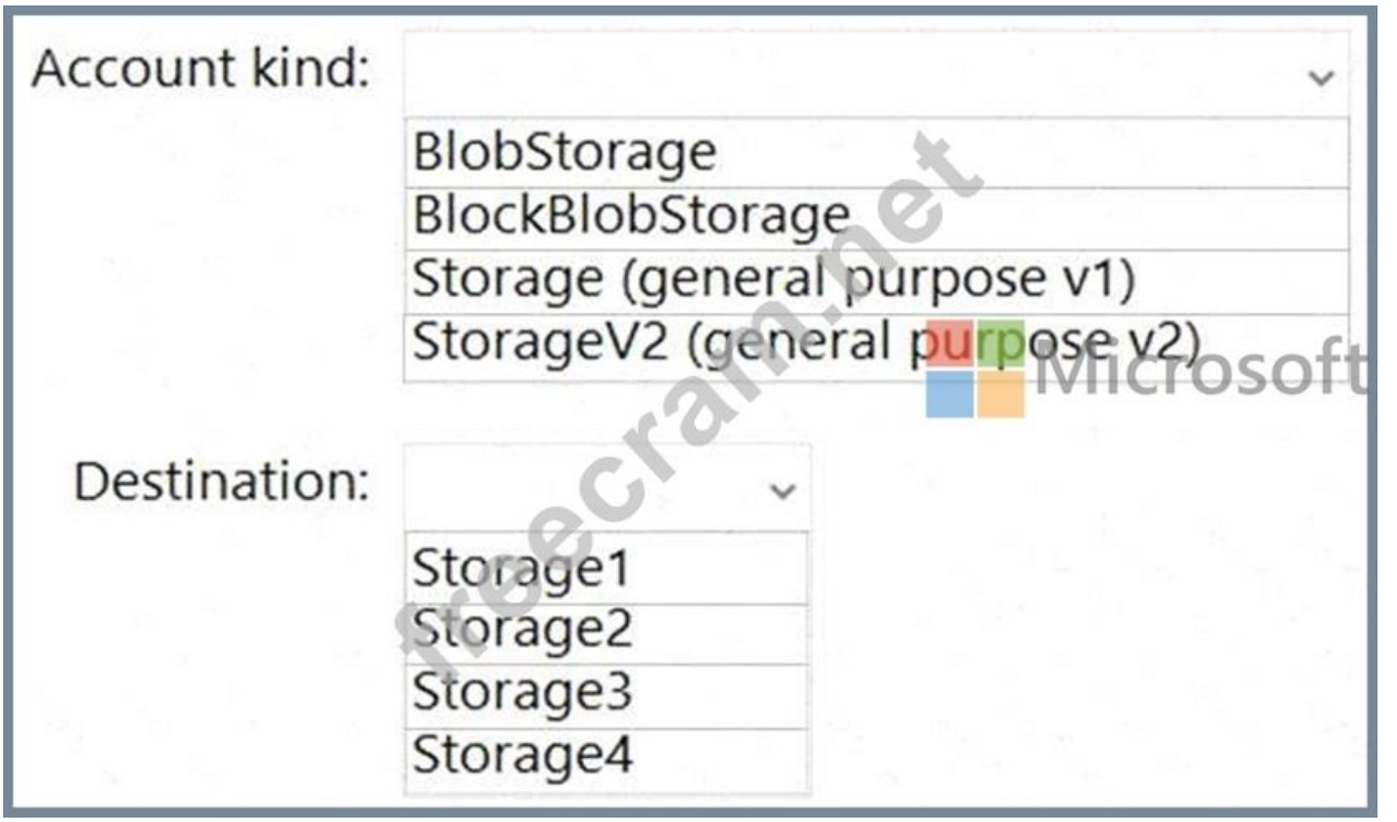

Answer:

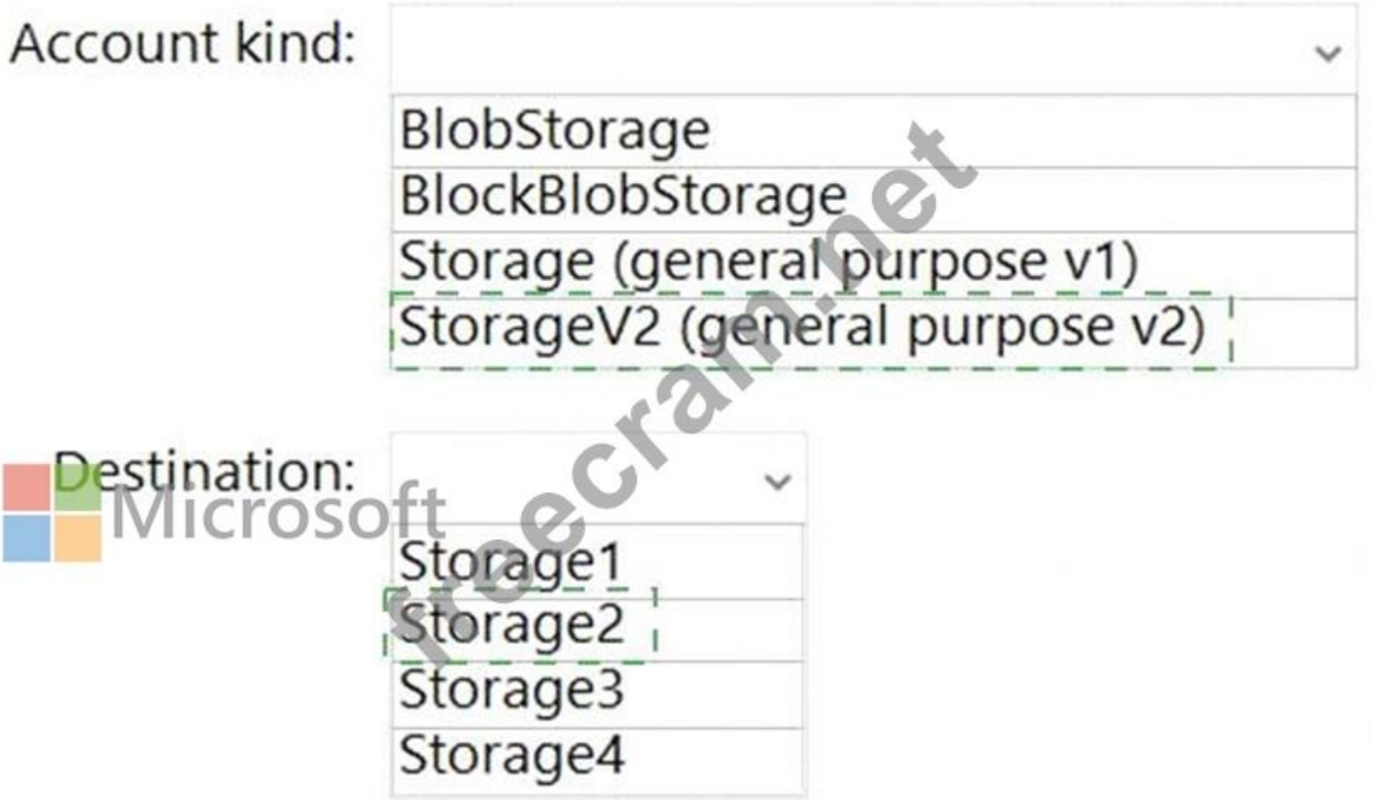

Explanation:

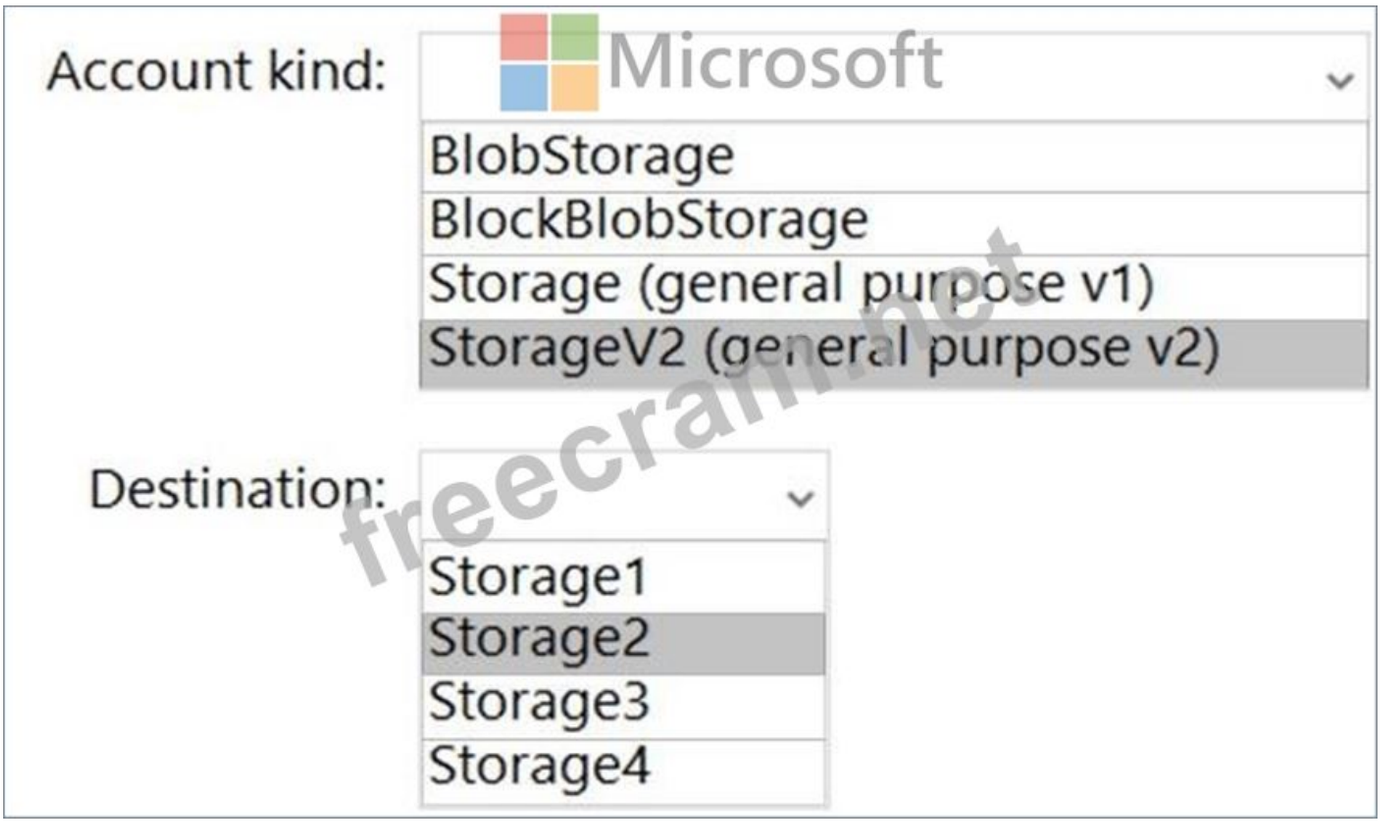

The question asks which storage account type should be created for storage5 and which existing storage account should be configured as the destination for replication - while supporting the planned changes.

Step 1: Planned Change Requirement

From the case study:

"Create a storage account named storage5 and configure storage replication for the Blob service." This indicates that storage5 will act as a secondary (destination) account to receive geo-replicated blobs (as part of RA-GRS or GRS replication configuration).

Therefore, the new storage5 must support Blob service replication features.

Step 2: Determining Correct Account Type

According to Microsoft Azure Administrator Documentation (Microsoft Learn: "Types of storage accounts"):

Account Type

Blob Support

File Share Support

Recommended Use

Storage (general purpose v1)

Basic blob

Limited

Legacy use, lacks advanced features

StorageV2 (general purpose v2)

Full blob, file, queue, table support

#

Supports replication (GRS, RA-GRS, ZRS) and access tiers

BlobStorage
Blob only
No file shares
Used only for block and append blobs
BlockBlobStorage
Blob only (premium tier)
No file shares
High-performance workloads only
Therefore, the correct type for storage5 must be StorageV2, as it supports both:
Blob service replication (GRS, RA-GRS, ZRS, LRS)
All tiers (Hot, Cool, Archive)
This meets the planned change and ensures the account supports replication for Blob service.
Step 3: Determining the Correct Destination Storage Account
The question also requires identifying which existing storage account should act as the destination for replication configuration.
In Azure, Blob service replication (for example, geo-redundant replication or object replication) requires:
Both storage accounts to be StorageV2.
Both to be in the same region pair and have object replication enabled.
Looking at the table:
Storage Account
Kind
Location
File Share
Identity-based Access
storage1
Storage (general purpose v1)
West US
sharea
Azure AD DS
storage2
StorageV2 (general purpose v2)
East US
shareb, sharec
Disabled
storage3
BlobStorage
East US 2
N/A
N/A
storage4

FileStorage
Central US
shared
Azure AD DS
Only storage2 meets all of the following conditions:
Type: StorageV2
Region: East US (a valid replication pair for East US 2 or secondary region)
Blob Service supported
Hence, storage2 is the correct destination storage account to configure for replication from storage5.
Step 4: Supporting Microsoft Azure Documentation Extracts
From Microsoft Learn - Azure Storage account overview:
"General-purpose v2 accounts support all Azure Storage features, including blob, file, queue, and table storage. They are recommended for most scenarios."
"Object replication requires both source and destination accounts to be of the same type (StorageV2) and in regions that are geo-paired."
"Blob and StorageV2 accounts support replication options such as LRS, ZRS, GRS, and RA-GRS."

NEW QUESTION: 112
You need to meet the connection requirements for the New York office.
What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

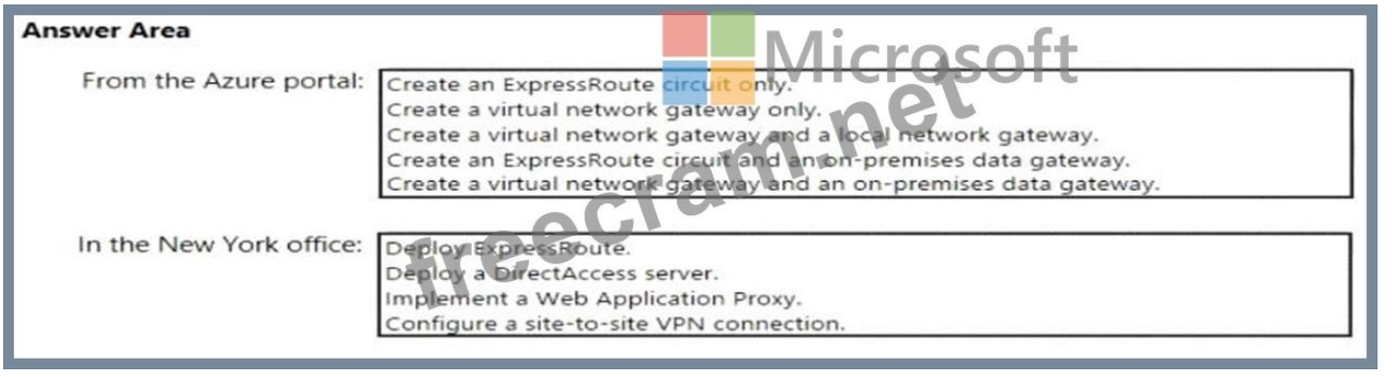

Answer:

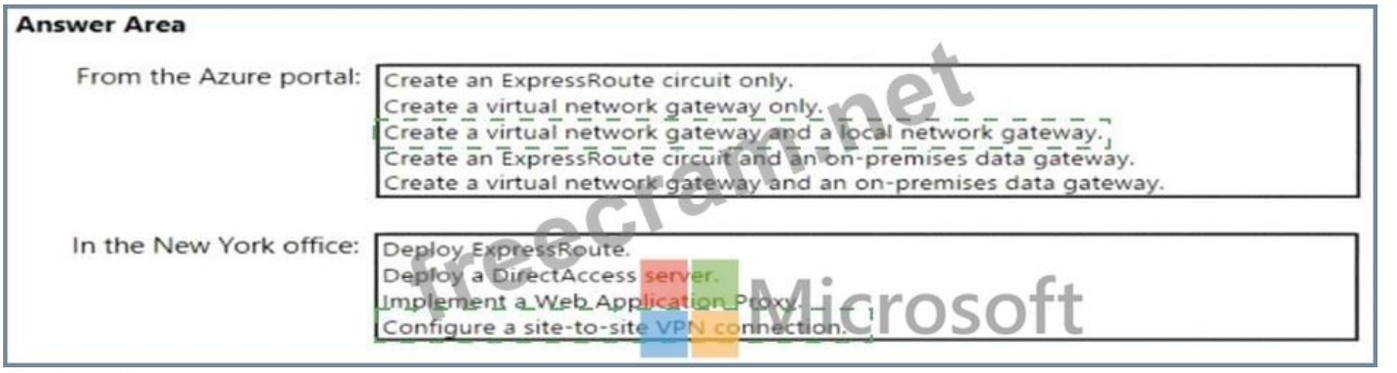

Explanation:

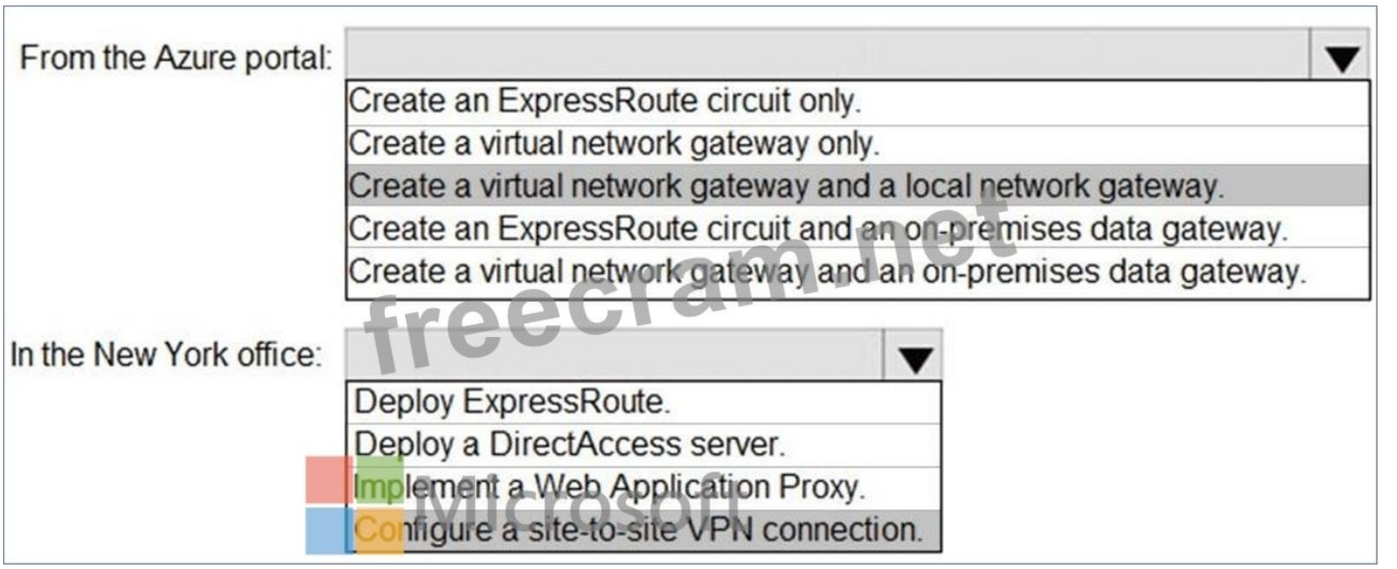

Box 1: Create a virtual network gateway and a local network gateway.

Azure VPN gateway. The VPN gateway service enables you to connect the VNet to the on-premises network through a VPN appliance. For more information, see Connect an on-premises network to a Microsoft Azure virtual network. The VPN gateway includes the following elements:

Virtual network gateway. A resource that provides a virtual VPN appliance for the VNet. It is responsible for routing traffic from the on-premises network to the VNet.

Local network gateway. An abstraction of the on-premises VPN appliance. Network traffic from the cloud application to the on-premises network is routed through this gateway.

Connection. The connection has properties that specify the connection type (IPSec) and the key shared with the on-premises VPN appliance to encrypt traffic.

Gateway subnet. The virtual network gateway is held in its own subnet, which is subject to various requirements, described in the Recommendations section below.

Box 2: Configure a site-to-site VPN connection

On premises create a site-to-site connection for the virtual network gateway and the local network gateway.

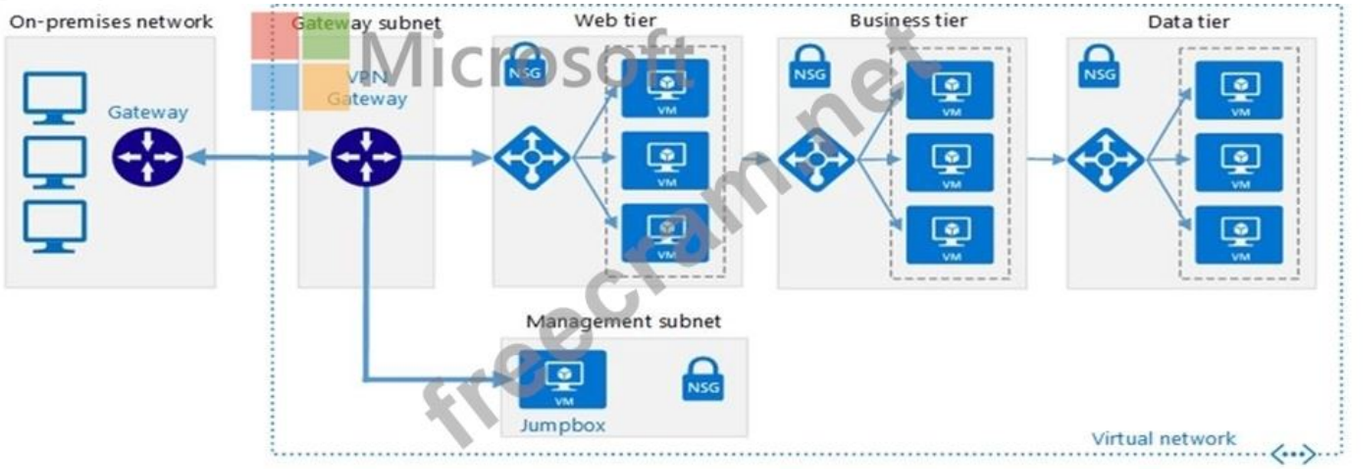

Scenario: Connect the New York office to VNet1 over the Internet by using an encrypted connection.

NEW QUESTION: 113

You have the Azure management groups shown in the following table.

Name	In management group
Tenant Root Group	*Not applicable*
ManagementGroup11	Tenant Root Group
ManagementGroup12	Tenant Root Group
ManagementGroup21	ManagementGroup11

You add Azure subscriptions to the management groups as shown in the following table.

Name	Management group
Subscription1	ManagementGroup21
Subscription2	ManagementGroup12

You create the Azure policies shown in the following table.

Name	Parameter	Scope
Not allowed resource types	virtualNetworks	Tenant Root Group
Allowed resource types	virtualNetworks	ManagementGroup12

For each of the following statements, select Yes it the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

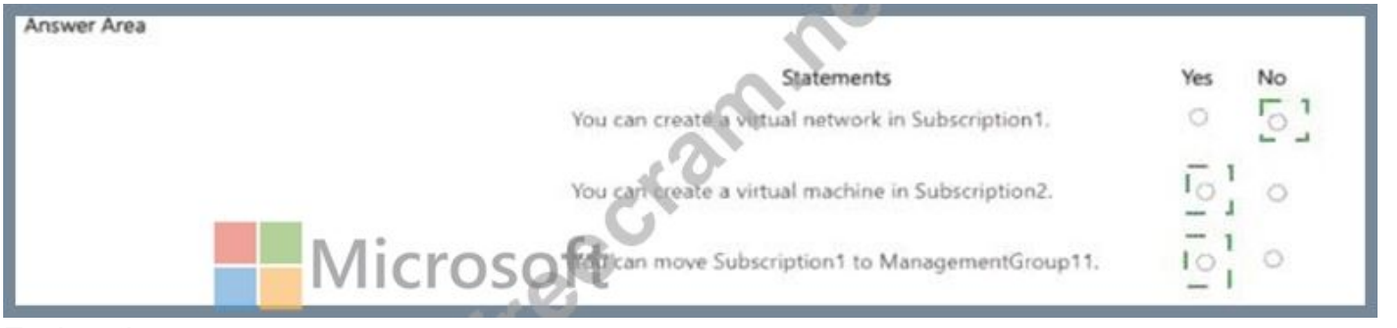

Explanation:

Azure management groups allow administrators to organize multiple subscriptions and apply governance policies such as Azure Policy and RBAC across a hierarchy. Policies assigned at a higher-level management group are inherited by all child management groups and subscriptions beneath it. When two or more policies conflict, Deny policies always override Allow policies as per Microsoft's official documentation (Microsoft Learn: Azure Policy Overview - Inheritance and Enforcement Logic).

In this question:
* Policy Assignments
* Tenant Root Group has a "Not allowed resource types" policy that denies virtualNetworks creation.
* ManagementGroup12 has an "Allowed resource types" policy that allows virtualNetworks creation.
Because Subscription2 resides under ManagementGroup12, it inherits the allowed policy.
However, Subscription1 resides under ManagementGroup21, which is a child of ManagementGroup11, both of which fall under the Tenant Root Group - inheriting the deny rule.
Conclusion:
* Subscription1 # Denied (inherits "Not allowed resource types: virtualNetworks")
* Subscription2 # Allowed (inherits "Allowed resource types: virtualNetworks")
* Creating Resources
* In Subscription1, you cannot create a virtual network due to the deny policy from the Tenant Root Group.
* In Subscription2, you can create a virtual machine, as no deny policy prevents it.
* Moving Subscriptions
* According to Azure governance hierarchy rules, subscriptions can be moved between management groups if the user has proper RBAC permissions (Owner or User Access Administrator). There are no policy restrictions preventing Subscription1 from being moved from ManagementGroup21 to ManagementGroup11, as this is within the same management group hierarchy.
Final Verified Answer (as per Microsoft Azure Administrator Documentation):
Statement
Answer
You can create a virtual network in Subscription1
No
You can create a virtual machine in Subscription2
Yes
You can move Subscription1 to ManagementGroup11
Yes
Official Microsoft Azure Reference (Document Extract Summary):
"Azure Policy effects are inherited by all child resources. A 'deny' effect from a parent management group overrides any 'allow' effects from descendant scopes. Policies at higher scopes take precedence in conflicts."
"Subscriptions can be moved between management groups within the same hierarchy when permissions are sufficient." (Source: Microsoft Learn - Azure Policy Overview, Management Groups Overview, and RBAC Permissions for Governance Management)
NEW QUESTION: 114
Your company purchases a new Azure subscription.
You create a file named Deploy json as shown in the following exhibit

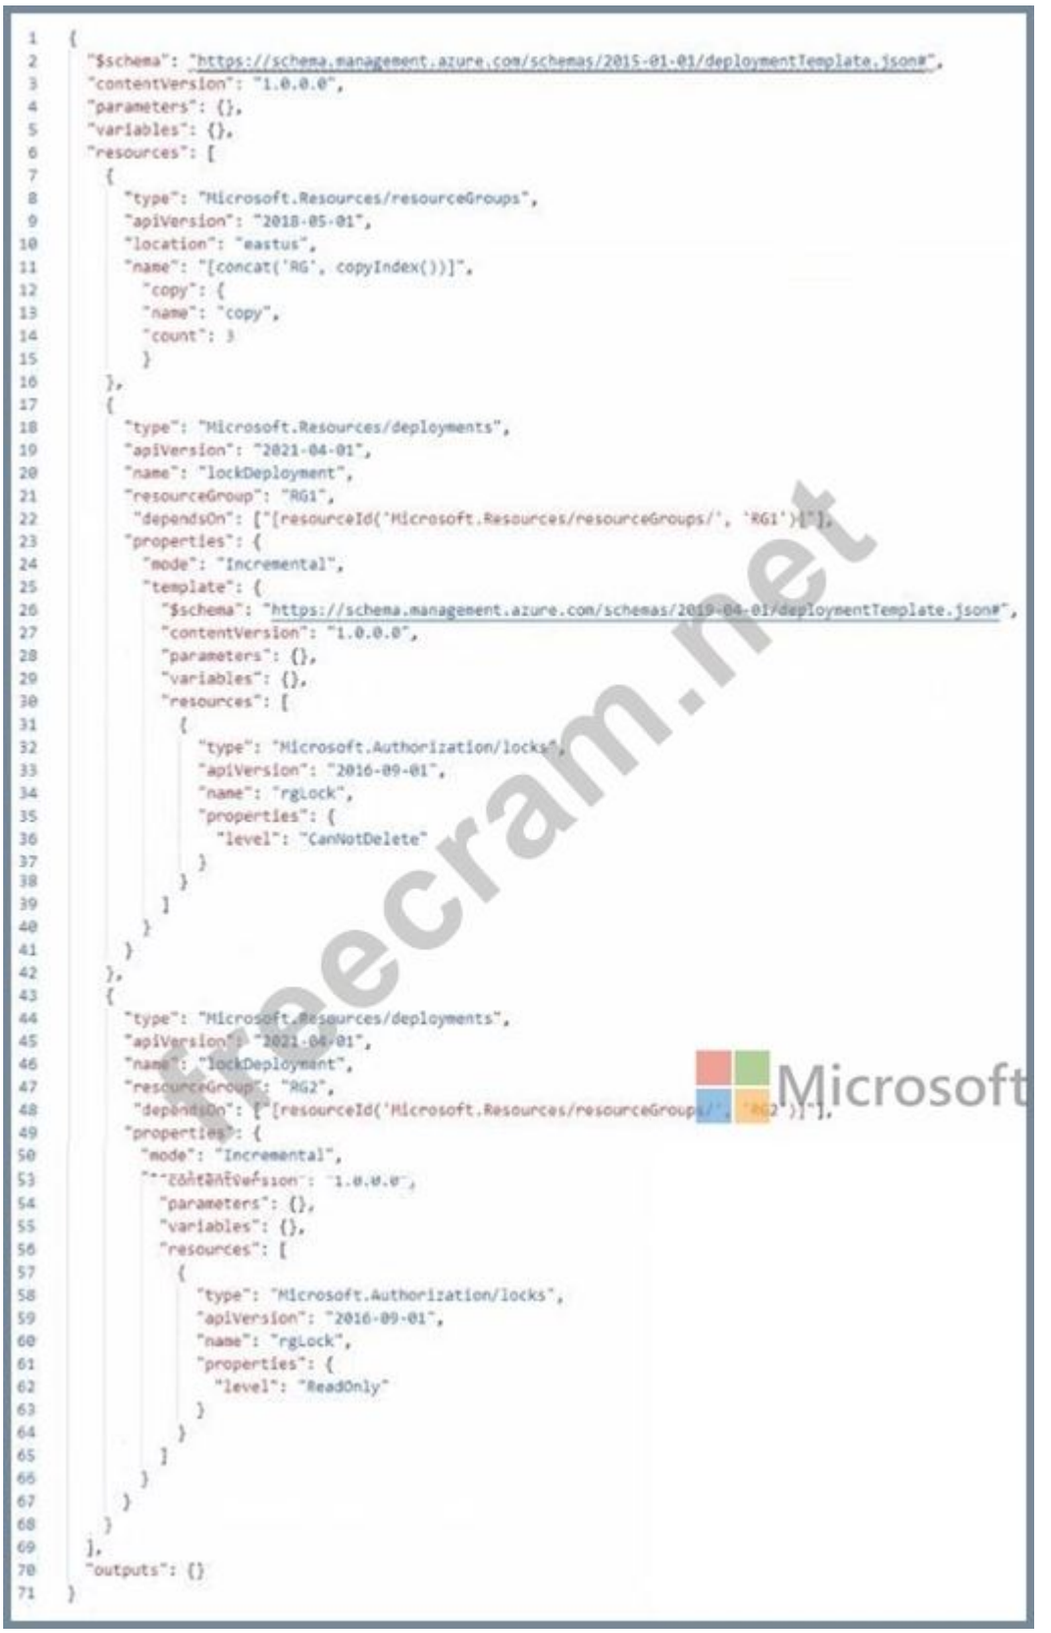

```
{
  "$schema": "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {},
  "variables": {},
  "resources": [
    {
      "type": "Microsoft.Resources/resourceGroups",
      "apiVersion": "2018-05-01",
      "location": "eastus",
      "name": "[concat('RG', copyIndex())]",
      "copy": {
      "name": "copy",
      "count": 3
      }
    },
    {
      "type": "Microsoft.Resources/deployments",
      "apiVersion": "2021-04-01",
      "name": "lockDeployment",
      "resourceGroup": "RG1",
      "dependsOn": ["[resourceId('Microsoft.Resources/resourceGroups/', 'RG1')]"],
      "properties": {
        "mode": "Incremental",
        "template": {
          "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
          "contentVersion": "1.0.0.0",
          "parameters": {},
          "variables": {},
          "resources": [
            {
              "type": "Microsoft.Authorization/locks",
              "apiVersion": "2016-09-01",
              "name": "rgLock",
              "properties": {
                "level": "CanNotDelete"
              }
            }
          ]
        }
      }
    },
    {
      "type": "Microsoft.Resources/deployments",
      "apiVersion": "2021-04-01",
      "name": "lockDeployment",
      "resourceGroup": "RG2",
      "dependsOn": ["[resourceId('Microsoft.Resources/resourceGroups/', 'RG2')]"],
      "properties": {
        "mode": "Incremental",
        "contentVersion": "1.0.0.0",
        "parameters": {},
        "variables": {},
        "resources": [
          {
            "type": "Microsoft.Authorization/locks",
            "apiVersion": "2016-09-01",
            "name": "rgLock",
            "properties": {
              "level": "ReadOnly"
            }
          }
        ]
      }
    }
  ],
  "outputs": {}
}
```


You connect to the subscription and run the following cmdlet:

New-AzDeployment -Location westus -TemplateFile "deploy.json""

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

Answer:

Answer Area

Explanation:

Based on the file named Deploy.json and the cmdlet you ran, here are the answers to your statements:

You can deploy a virtual machine to RGI. = No

You can deploy a virtual machine to RG2. = No

You can manually create a resource group named RG3. = Yes

Let me explain why:

The Deploy.json file defines a template for creating a resource group and a virtual machine in Azure. The template has two parameters: resourceGroupName and vmName. The template also has two resources: one for the resource group and one for the virtual machine. The resource group resource has a property called name, which is set to the value of the resourceGroupName parameter. The virtual machine resource has a property called location, which is set to the value of the location parameter of the deployment cmdlet.

The cmdlet you ran specifies the location as westus and the template file as Deploy.json. However, it does not specify any values for the resourceGroupName and vmName parameters. Therefore, the cmdlet will prompt you to enter those values interactively before creating the deployment.

If you enter RGI as the value for the resourceGroupName parameter and VM1 as the value for the vmName parameter, then the cmdlet will create a resource group named RGI and a virtual machine named VM1 in the westus location. Therefore, you can deploy a virtual machine to RGI.

However, if you enter RG2 as the value for the resourceGroupName parameter, then the cmdlet will fail with an error. This is because RG2 already exists in your subscription and you cannot create a resource group with the same name as an existing one. Therefore, you cannot deploy a virtual machine to RG2 using this template and cmdlet.

You can manually create a resource group named RG3 by using another cmdlet: New-AzResourceGroup.

This cmdlet takes two parameters: Name and Location. For example, you can run the following

cmdlet to create a resource group named RG3 in westus:
New-AzResourceGroup -Name RG3 -Location westus

NEW QUESTION: 115
You have an Azure subscription that contains the resources shown in the following table.

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

You need to load balance HTTPS connections to vm1 and vm2 by using lb1.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

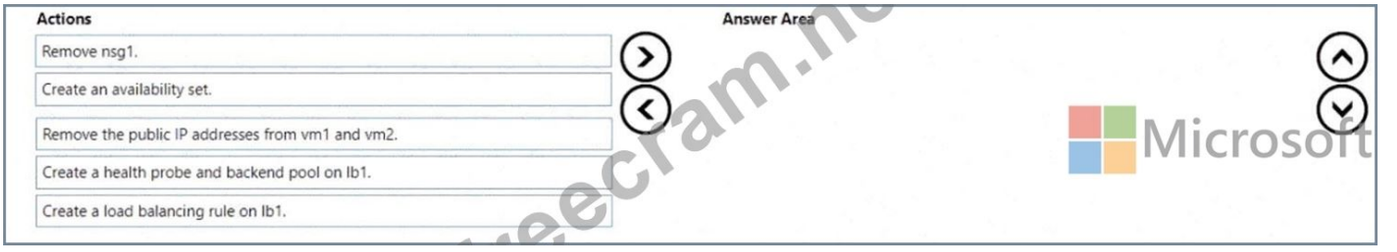

Answer:

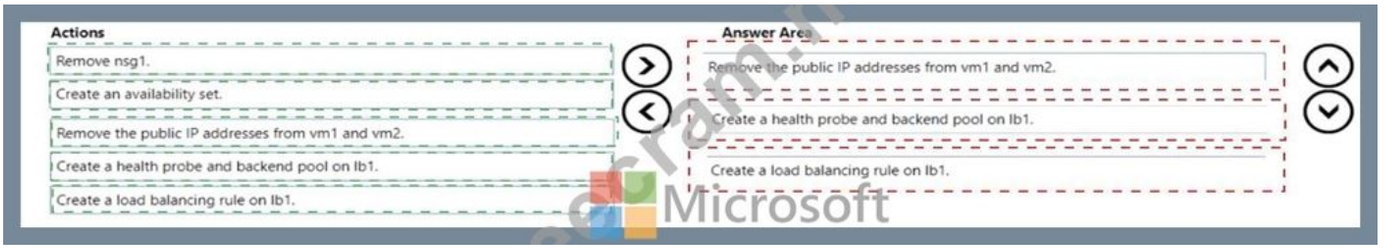

Explanation:

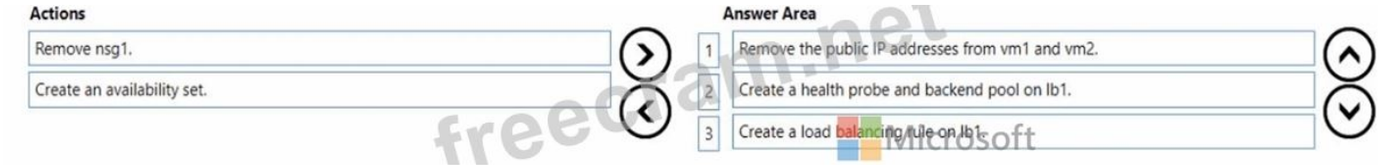

To configure an Azure Standard Load Balancer to distribute HTTPS (TCP port 443) traffic between two virtual machines (vm1 and vm2), several specific configuration steps must be followed according to Microsoft Azure Administrator documentation ("Load Balancer Standard SKU - configuration and differences").
* 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: 116
You have two Azure App Service apps named App1 and App2. Each app has a production deployment slot and a test deployment slot. The Backup Configuration settings for the production slots are shown in the following table.

App	Backup Every	Start backup schedule from	Retention (Days)	Keep at least one backup
App1	1 Days	January 6, 2021	0	Yes
App2	1 Days	January 6, 2021	30	Yes

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer:

Explanation:

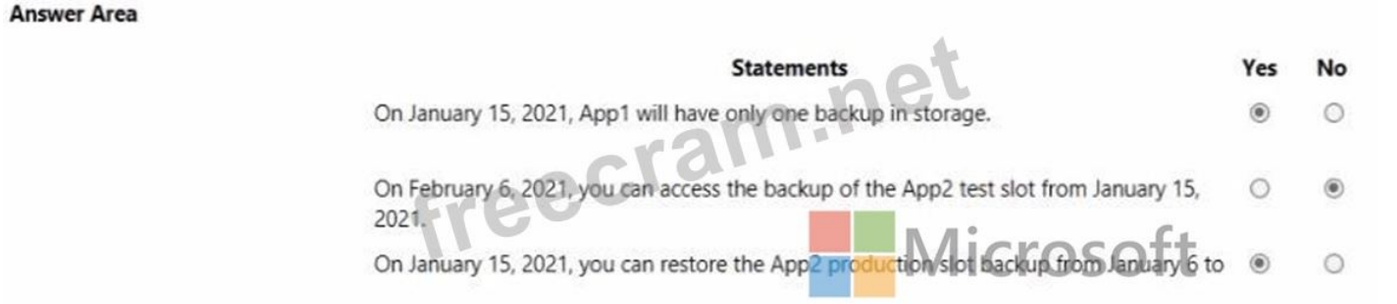

On January 15, 2021, App1 will have only one backup in storage. Yes, this is correct. According to the table, App1 has a backup every 1 day, starting from January 6, 2021, with a retention of 0

days. This means that each backup will be deleted after 0 days, or as soon as the next backup is created. Therefore, on January 15,
2021, App1 will have only one backup in storage, which is the one created on that day1.
On February 6, 2021, you can access the backup of the App2 test slot from January 15, 2021.
No, this is not correct. According to the table, App2 has a backup every 1 day, starting from January 6, 2021, with a retention of 30 days. This means that each backup will be deleted after 30 days, or when the storage limit is reached. However, the table also shows that App2 has a setting of "Keep at least one backup" set to Yes. This means that the oldest backup will be retained even if it exceeds the retention period or the storage limit2.
Therefore, on February 6, 2021, you can access the backup of the App2 test slot from January 6, 2021, but not from January 15, 2021.
On January 15, 2021, you can restore the App2 production slot backup from January 6 to the App2 test slot. Yes, this is correct. According to the web search results, you can restore a backup by overwriting an existing app or by restoring to a new app or slot3. You can also restore a backup from a different slot or app as long as they are in the same subscription and region4.
Therefore, on January 15, 2021, you can restore the App2 production slot backup from January 6 to the App2 test slot.

NEW QUESTION: 117
You need to implement the planned changes for the storage account content. Which containers and file shares can you use to organize the content?
A. share1 only
B. cont1 and share1 only
C. share1 and share2 only
D. cont1, share1, and share2 only
E. cont1, cont2, share1, and share2
Answer: (SHOW ANSWER)
In the scenario, storage1 is configured as StorageV2 with Hierarchical namespace = Yes, while storage2 is configured as StorageV2 with Hierarchical namespace = No.
From Microsoft's Azure Storage Documentation and AZ-104 Study Guide, the following principles apply:
* A hierarchical namespace (enabled when the storage account has Azure Data Lake Storage Gen2 capabilities) allows the use of directories within containers to organize data.
* The hierarchical namespace provides directory and file-level structure similar to a file system. This is supported only for blob containers, not for Azure Files.
* Azure Files (file shares) do not depend on hierarchical namespaces and cannot have directories in the same way Data Lake Gen2 does - directories can exist inside the share but not in the blob container sense.
* The planned change states that you must use directories whenever possible to organize content.
Therefore, only storage accounts with hierarchical namespace enabled can use directory

structures - that's storage1.
In this case:
* storage1 (Hierarchical namespace = Yes) # supports containers (like cont1) and file shares (like share1).
* storage2 (Hierarchical namespace = No) # does not support directories within blob containers (Data Lake structure).
Hence, you can use only cont1 (container in storage1) and share1 (file share in storage1) to organize content as required.
This is directly supported by the Microsoft documentation on Data Lake Storage Gen2:
"When you enable the hierarchical namespace for a storage account, you can organize objects into directories and subdirectories. This capability is available only for accounts configured for Data Lake Storage Gen2." Final Verified Answer: # B. cont1 and share1 only

NEW QUESTION: 118
Your network contains an on-premises Active Directory Domain Services (AD DS) domain.
The domain contains the identities shown in the following table.

Name	Description	In organizational unit (OU)
User1	User	OU2
User2	User	OU1
Group1	Global group that contains User1	OU1

You have an Azure subscription that uses Microsoft Entra Connect sync to sync OU1 from the AD DS domain to the Microsoft Entra tenant.
The Microsoft Entra tenant contains a cloud only user named User3.

Answer:

Explanation:
NO
YES
NO
In this scenario, Microsoft Entra Connect (formerly Azure AD Connect) is configured to synchronize only the Organizational Unit (OU1) from the on-premises Active Directory Domain

Services (AD DS) environment to the Microsoft Entra tenant (Azure AD).
The synchronization scope determines which objects (users, groups, devices) are replicated to Azure AD. Any objects outside the selected OUs are not synchronized and therefore do not exist in Azure AD.
Here's what happens to each user:
* User1 (in OU2) - Because OU2 is not included in the synchronization scope, User1 is not synchronized to Azure AD. Therefore, User1 does not have an account in Microsoft Entra ID and cannot authenticate to any cloud-based resource such as Share1 (which is assumed to be an Azure file share or Microsoft
365 resource).
* Result: No access.
* User2 (in OU1) - Since OU1 is synchronized using Microsoft Entra Connect, User2 exists in Azure AD. User2's identity is recognized by Azure AD and can be used to access cloud-based resources (like Share1) if permissions are assigned appropriately (either directly or via a group such as Group1).
* Result: Yes, access is possible.
* User3 (cloud-only account) - While User3 exists in Azure AD, there's no indication that this user has been assigned access to Share1. Because access is granted through synchronization (from AD DS) and group membership in synchronized objects (like Group1), User3 would not automatically inherit those permissions.
* Result: No access.
Additionally, Group1, which contains User1 and resides in OU1, will be synchronized. However, since User1 is in OU2, User1's membership in Group1 is not synchronized, because non-synchronized objects cannot appear in synchronized group memberships in Azure AD.
This behavior is defined by Microsoft's Azure AD Connect synchronization rules, which state: "Objects outside of the configured synchronization scope are not synchronized to Azure AD. Group memberships including users from outside the synchronization scope will not include those users in Azure AD."

NEW QUESTION: 119
You need to resolve the licensing issue before you attempt to assign the license again.
What should you do?
A. From the Groups blade, invite the user accounts to a new group.
B. From the Profile blade, modify the usage location.
C. From the Directory role blade, modify the directory role.
Answer: (SHOW ANSWER)
Scenario: Licensing Issue
1. You attempt to assign a license in Azure to several users and receive the following error message:
"Licenses not assigned. License agreement failed for one user."
2. You verify that the Azure subscription has the available licenses.

Solution:
License cannot be assigned to a user without a usage location specified.
Some Microsoft services aren't available in all locations because of local laws and regulations.
Before you can assign a license to a user, you must specify the Usage location property for the user. You can specify the location under the User > Profile > Settings section in the Azure portal.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/users-groups-roles/licensing-groups-resolve-problems

NEW QUESTION: 120
You have a Microsoft Entra tenant that contains the groups shown in the following table.

Name	Type	Has an assigned license
Group1	Security	Yes
Group2	Security	No
Group3	Microsoft 365	Yes
Group4	Microsoft 365	No

The tenant contains the users shown in the following table.

Name	Member of	Has a direct assigned license
User1	None	Yes
User2	Group1	No
User3	Group4	Yes
User4	None	No

Which users and groups can you delete? To answer, select the appropriate options in the answer area.
NOTE Each correct selection is worth one point.
Answer Area

Answer:

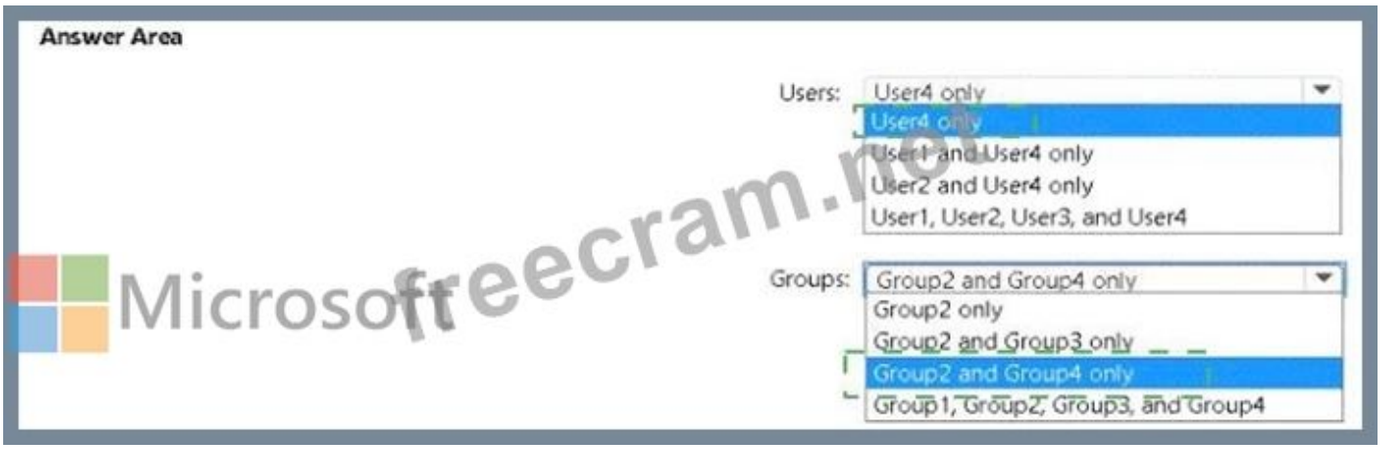

Explanation:

NEW QUESTION: 121

You have an Azure subscription that contains an Azure Stream Analytics job named Job1.

You need to monitor input events for Job1 to identify the number of events that were NOT processed.

Which metric should you use?

A. Output Events

B. Backlogged Input Events

C. Out-of-Order Events

D. Late Input Events

Answer: (SHOW ANSWER)

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: 122

You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Resource group	Location
RG1	Resource group	Not applicable	Central US
RG2	Resource group	Not applicable	West US
VMSS1	Virtual machine scale set	RG2	West US
Proximity1	Proximity placement group	RG1	West US
Proximity2	Proximity placement group	RG2	Central US
Proximity3	Proximity placement group	RG1	Central US

You need to configure a proximity placement group for VMSS1.

Which proximity placement groups should you use?

A. Proximity2 only

B. Proximity 1, Proximity2, and Proximity3

C. Proximity 1 and Proximity3 only

D. Proximity1 only

Answer: (SHOW ANSWER)

Placement Groups is a capability to achieve co-location of your Azure Infrastructure as a Service (IaaS) resources and low network latency among them, for improved application performance. Azure proximity placement groups represent a new logical grouping capability for your Azure Virtual Machines, which in turn is used as a deployment constraint when selecting where to place your virtual machines. In fact, when you assign your virtual machines to a proximity placement group, the virtual machines are placed in the same data center, resulting in lower and deterministic latency for your applications.

The VMSS should share the same region, even it should be the same zone as proximity groups are located in the same data center. Accordingly, it should be proximity 2 only.

Reference:

https://azure.microsoft.com/en-us/blog/introducing-proximity-placement-groups

Name	Kind	Redundancy
storage1	StorageV2	Geo-zone-redundant storage (GZRS)
storage2	BlobStorage	Read-access geo-redundant storage (RA-GRS)
storage3	BlockBlobStorage	Zone-redundant storage (ZRS)

following table.

You need to identify which storage accounts support lifecycle management, and which storage accounts support moving data to the Archive access tier. What should you identify for each

requirement? To answer, select the appropriate options in the answer area. NOTE: Each correct answer is worth one point.

Answer:

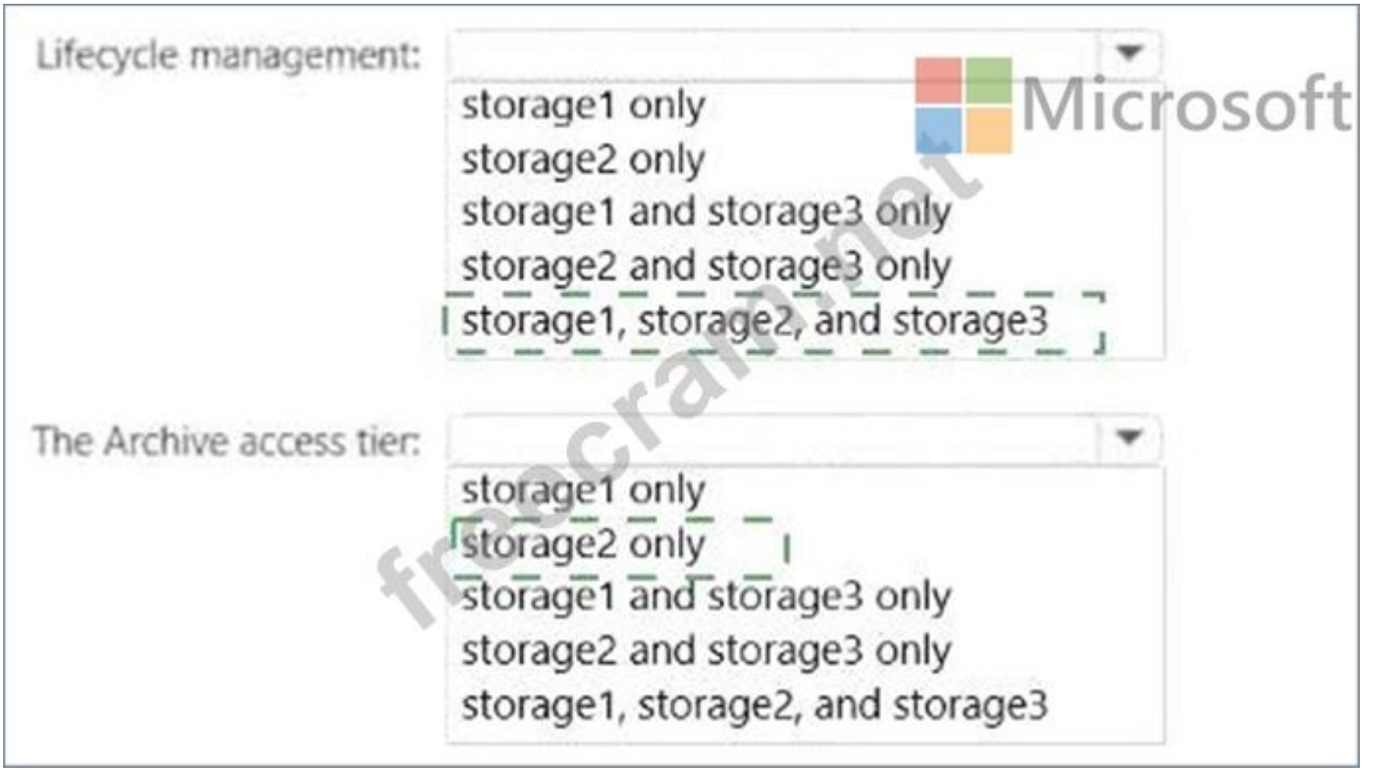

Explanation:

1 - storage1, storage2, storage3

2 - storage2

NEW QUESTION: 124

You need to create container1 and share1.

Which storage accounts should you use for each resource? To answer, select the appropriate options in t he answer area.

NOTE: Each correct selection is worth one point.

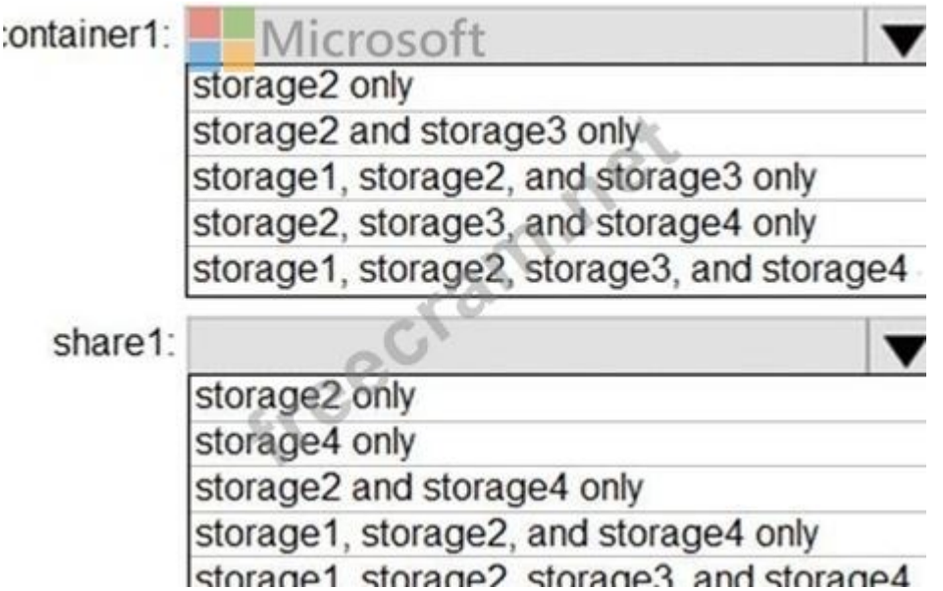

Answer:

In Azure, the type of storage account determines which services are available and which features can be configured. Let's analyze the given storage accounts in this case study:

Storage Account
Kind
Services Supported
Hierarchical Namespace
File Share Access (Identity-based)
storage1
General-purpose v1
Blobs, Files, Queues, Tables
No
Azure AD DS
storage2
General-purpose v2
Blobs, Files, Queues, Tables
No
Disabled
storage3
BlobStorage
Blobs only
N/A
N/A
storage4
FileStorage
Files only
Azure AD DS

Now, based on Azure Administrator (AZ-104) documentation:

Blob containers can only be created in StorageV2 (general-purpose v2) or BlobStorage accounts.

Azure file shares can only be created in StorageV2 or FileStorage accounts.
General-purpose v1 (StorageV1) accounts are considered legacy and should not be used for new resource creation when a StorageV2 or Blob/File-specific account is available.
FileStorage accounts are specifically designed for premium performance file shares (Azure Files Premium tier).
Identity-based access to file shares using Azure Active Directory Domain Services (Azure AD DS) is only supported on StorageV2 and FileStorage types with Azure AD DS integration enabled.
Step-by-step reasoning:
For container1 (Blob container)
The blob service requires either a StorageV2 or BlobStorage type.
storage2 (StorageV2) and storage3 (BlobStorage) both meet this requirement.
storage1 (StorageV1) is not recommended for new blob container creation due to lack of modern capabilities such as access tiers and lifecycle management.# Therefore, container1 should be created in storage2 and storage3 only.
For share1 (File share)
Azure file shares can only exist in StorageV2 or FileStorage accounts.
From the table:
storage2 (StorageV2) supports file shares (shareb, sharec).
storage4 (FileStorage) supports premium Azure Files.
storage1 (StorageV1) supports file shares but is legacy and not recommended for new configurations.
storage3 (BlobStorage) cannot host file shares at all.# Therefore, share1 should be created in storage2 and storage4 only.
Final Verified Answer:# container1 # storage2 and storage3 only# share1 # storage2 and storage4 only Source:
Microsoft Learn: Azure Storage account overview
Microsoft Learn: Create a storage account to store blobs, files, queues, and tables Microsoft Learn: Identity-based access to Azure Files using Azure AD DS Microsoft Learn: Azure Blob storage tiers and account types

NEW QUESTION: 125
You have an Azure subscription that contains the virtual machines shown in the following table.
You deploy a load balancer that has the following configurations:
* Name: LB 1
* Type: Internal
* SKU: Standard
* Virtual network: VNET1
You need to ensure that you can add VM1 and VM2 to the backend pool of L81.
Solution: You create two Standard SKU public IP addresses and associate a Standard SKU public IP address to the network interface of each virtual machine.
Does this meet the goal?

A. Yes
B. No
Answer: (SHOW ANSWER)
In Azure, load balancer configurations must have consistent SKU compatibility between the load balancer and its associated backend pool members. The Standard SKU Load Balancer requires that all backend pool members - such as virtual machines, network interfaces, or virtual machine scale sets - use Standard SKU public IPs or Standard SKU configurations if external, and proper network interface associations if internal.
However, in this question:
* LB1 is an internal Standard SKU load balancer.
* Internal load balancers do not require or support public IP addresses on the backend virtual machines.
* Backend pool members of an internal load balancer are identified by private IP addresses from the virtual network, not public IPs.
According to Microsoft Azure Load Balancer documentation (Azure Administrator Guide, "Configure Standard Load Balancer"),
"For a Standard internal load balancer, the backend pool can contain only virtual machines or NICs that are in the same virtual network as the load balancer. Public IP addresses are not required and not supported for backend members." To correctly meet the goal, the proper solution is to:
* Ensure that VM1 and VM2 are deployed within the same virtual network (VNET1) and subnet as LB1.
* Ensure that the NICs of both VMs are configured with Standard SKU (if public IPs are attached elsewhere) to match the load balancer SKU.
* Add these NICs or virtual machines directly to the backend pool of LB1.
Creating and assigning Standard SKU public IP addresses to each VM is unnecessary and does not help connect them to the internal load balancer. In fact, it contradicts Azure's design for internal load balancers, which only route traffic using private addresses within the virtual network scope.
Therefore, the given solution does not meet the goal.

NEW QUESTION: 126
You need to prepare the environment to ensure that the web administrators can deploy the web apps as quickly as possible.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

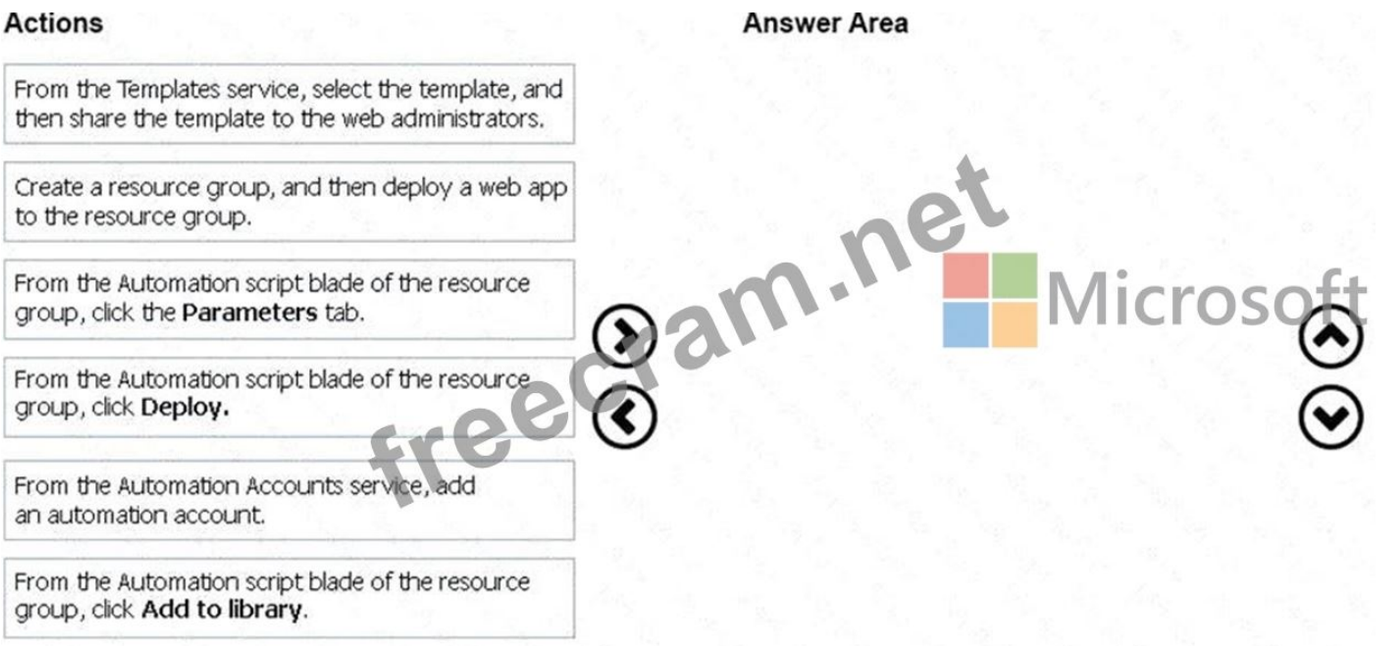

Answer:

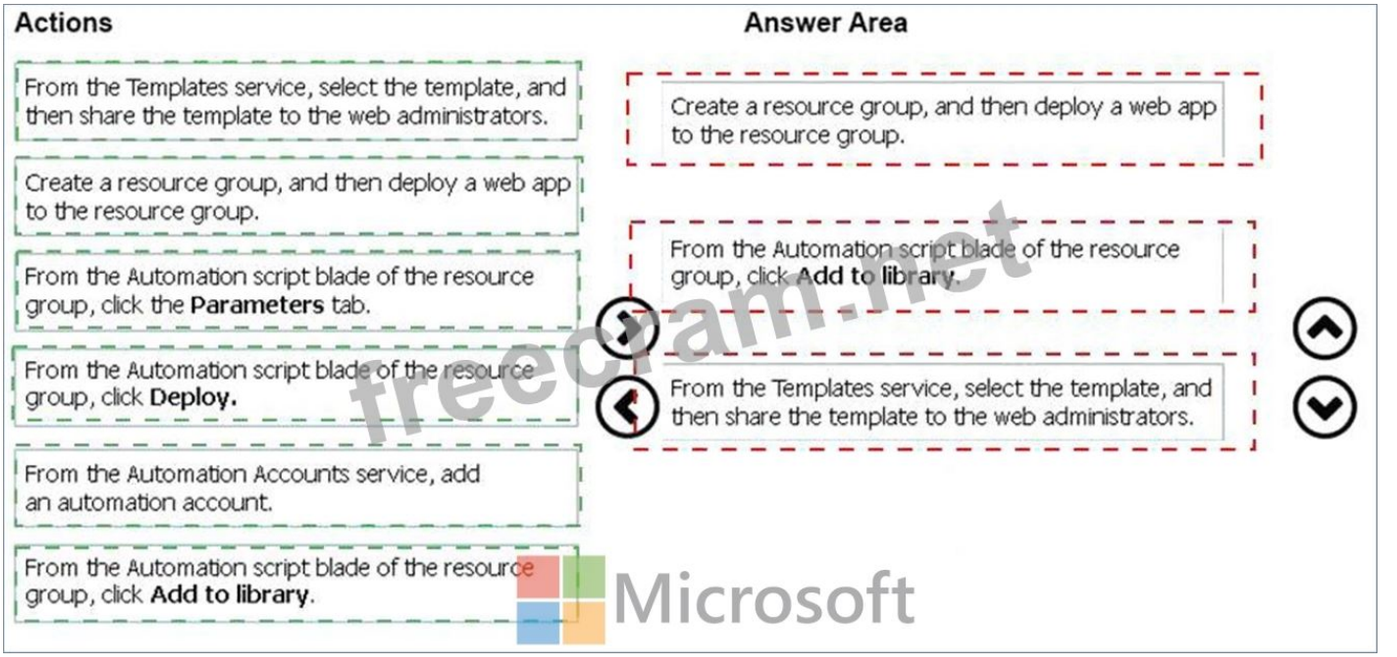

Explanation:

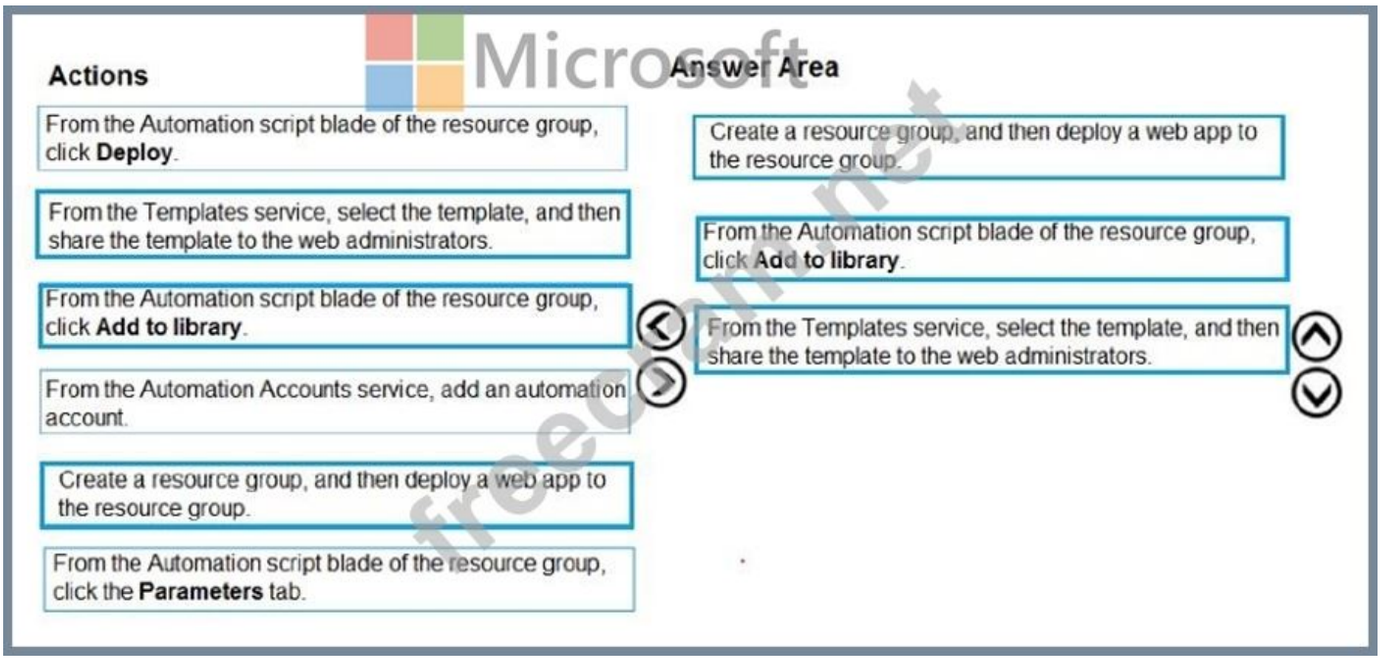

Scenario:

1. Web administrators will deploy Azure web apps for the marketing department.
2. Each web app will be added to a separate resource group.
3. The initial configuration of the web apps will be identical.
4. The web administrators have permission to deploy web apps to resource groups.

Steps:

1 --> Create a resource group, and then deploy a web app to the resource group.

2 --> From the Automation script blade of the resource group , click Add to Library.

3 --> From the Templates service, select the template, and then share the template to the web administrators .

References:

https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/quickstart-create-templates-use-the- portal

NEW QUESTION: 127

You have an Azure Active Directory (Azure AD) tenant.

You plan to delete multiple users by using Bulk delete in the Azure Active Directory admin center.

You need to create and upload a file for the bulk delete.

Which user attributes should you include in the file?

A. The user principal name and usage location of each user only

B. The user principal name of each user only

C. The display name of each user only

D. The display name and usage location of each user only

E. The display name and user principal name of each user only

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: 128
You have a registered DNS domain named contoso.com.
You create a public Azure DNS zone named contoso.com.
You need to ensure that records created in the contoso.com zone are resolvable from the internet.
What should you do?
A. Modify the NS records in the DNS domain registrar.
B. Create NS records in contoso.com.
C. Create the SOA record in contoso.com.
D. Modify the SOA record in the DNS domain registrar.
Answer: (SHOW ANSWER)
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: 129
You have an Azure AD tenant named adatum.com that contains the groups shown in the following table.

Name	Member of
Group1	*None*
Group2	Group1
Group3	Group2

Adatum.com contains the users shown in the following table.

Name	Member of
User1	Group1
User2	Group2
User3	Group3
User4	*None*

You assign the Azure AD Premium P2 license to Group I and User4.
Which users are assigned the Azure AD Premium P2 license?
A. User4 only
B. User1 and User4 only
C. User1. User2. and User4 only
D. User1, User2, User3, and User4
Answer: (SHOW ANSWER)
According to the Microsoft documentation, when you assign a license to a group, all members of that group are automatically assigned the license. However, if a user is already assigned the same license directly or through another group, the license is not duplicated.
In your scenario, you assigned the Azure AD Premium P2 license to Group1 and User4. This means that all members of Group1, which are User1 and User2, will also get the license. User4 will get the license directly.
User3 will not get the license because they are not a member of Group1 or assigned the license directly.
Therefore, the users who are assigned the Azure AD Premium P2 license are User1, User2, and User4 only.

NEW QUESTION: 130
You need to resolve the Active Directory issue.
What should you do?

A. From Active Directory Users and Computers, select the user accounts, and then modify the User Principal Name value.
B. Run idfix.exe, and then use the Edit action.
C. From Active Directory Domains and Trusts, modify the list of UPN suffixes.
D. From Azure AD Connect, modify the outbound synchronization rule.
Answer: B (LEAVE A REPLY)
IdFix is used to perform discovery and remediation of identity objects and their attributes in an on-premises Active Directory environment in preparation for migration to Azure Active Directory.
IdFix is intended for the Active Directory administrators responsible for directory synchronization with Azure Active Directory.
Scenario: Active Directory Issue
Several users in humongousinsurance.com have UPNs that contain special characters.
You suspect that some of the characters are unsupported in Azure AD.
References: https://www.microsoft.com/en-us/download/details.aspx?id=36832

NEW QUESTION: 131
You plan to create an Azure Storage account named storage1 that will contain a file share named share1.
You need to ensure that share! can support SMB Multichannel. The solution must minimize costs.
How should you configure storage1?
A. Standard performance with locally-redundant storage (IRS)

B. Premium performance with locally-redundant storage (LRS)
C. Standard performance with zone-redundant storage (ZRS)
Answer: (SHOW ANSWER)
SMB Multichannel is a feature of Azure Files that enables clients to establish multiple network connections to an Azure file share simultaneously. This allows parallel data transfer, thereby increasing throughput and providing fault tolerance for high-performance workloads.
According to the Microsoft Azure Administrator documentation and Azure Files technical reference, SMB Multichannel is available only for premium file shares hosted on Azure Files Premium Storage accounts that use locally redundant storage (LRS) replication.
Extract from Microsoft documentation:
"SMB Multichannel is supported only for premium file shares. Premium file shares are hosted in FileStorage storage accounts, which use SSD-based performance and locally redundant storage (LRS). Standard (HDD- based) file shares do not support SMB Multichannel." The Premium tier uses SSD storage, designed for low latency and high IOPS workloads, and is required for enabling SMB Multichannel. Using LRS (Locally Redundant Storage) provides the lowest cost redundancy option while still ensuring three replicas within the same datacenter.
Therefore, to meet the requirement of SMB Multichannel while minimizing costs, you must choose Premium performance with LRS redundancy.

NEW QUESTION: 132
You need to configure Azure Backup to back up the file shares and virtual machines.
What is the minimum number of Recovery Services vaults and backup policies you should create? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Recovery Services vaults

1
2
3
4
7

Backup policies

1
2
3
4
5
6

Answer:

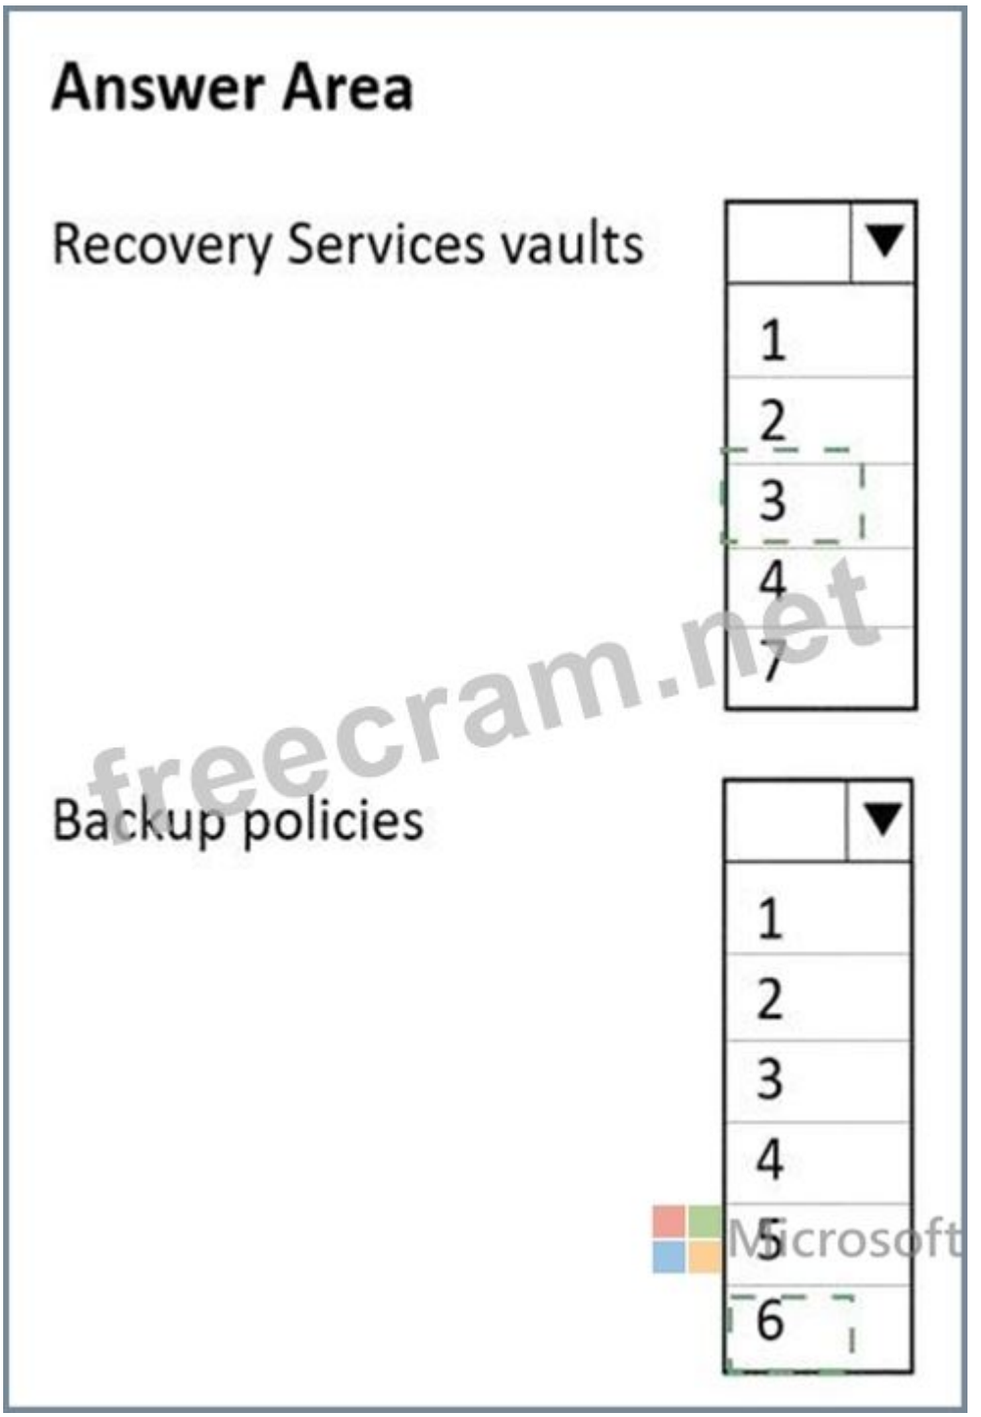

Explanation:

NEW QUESTION: 133

You need to prepare the environment to meet the authentication requirements.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Allow inbound TCP port 8080 to the domain controllers in the Miami office.
B. Add
http://autogon.microsoftazuread-sso.com to the intranet zone of each client computer in the Miami office.
C. Join the client computers in the Miami office to Azure AD.
D. Install the Active Directory Federation Services (AD FS) role on a domain controller in the Miami office.
E. Install Azure AD Connect on a server in the Miami office and enable Pass-through Authentication.
Answer: (SHOW ANSWER)
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
E: Seamless SSO works with any method of cloud authentication - Password Hash Synchronization or Pass- through Authentication, and can be enabled via Azure AD Connect.
References:
https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-sso-quick-start

NEW QUESTION: 134
You have a deployment template named Template1 that is used to deploy 10 Azure web apps.
You need to identify what to deploy before you deploy Tempi ate 1. The solution must minimize Azure costs, What should you identify?
A. five Azure Application Gateways
B. one App Service plan
C. 10 App Service plans
D. one Azure Traffic Manager
E. one Azure Application Gateway
Answer: (SHOW ANSWER)

NEW QUESTION: 135
You have an Azure subscription that contains the storage accounts shown in the following exhibit.

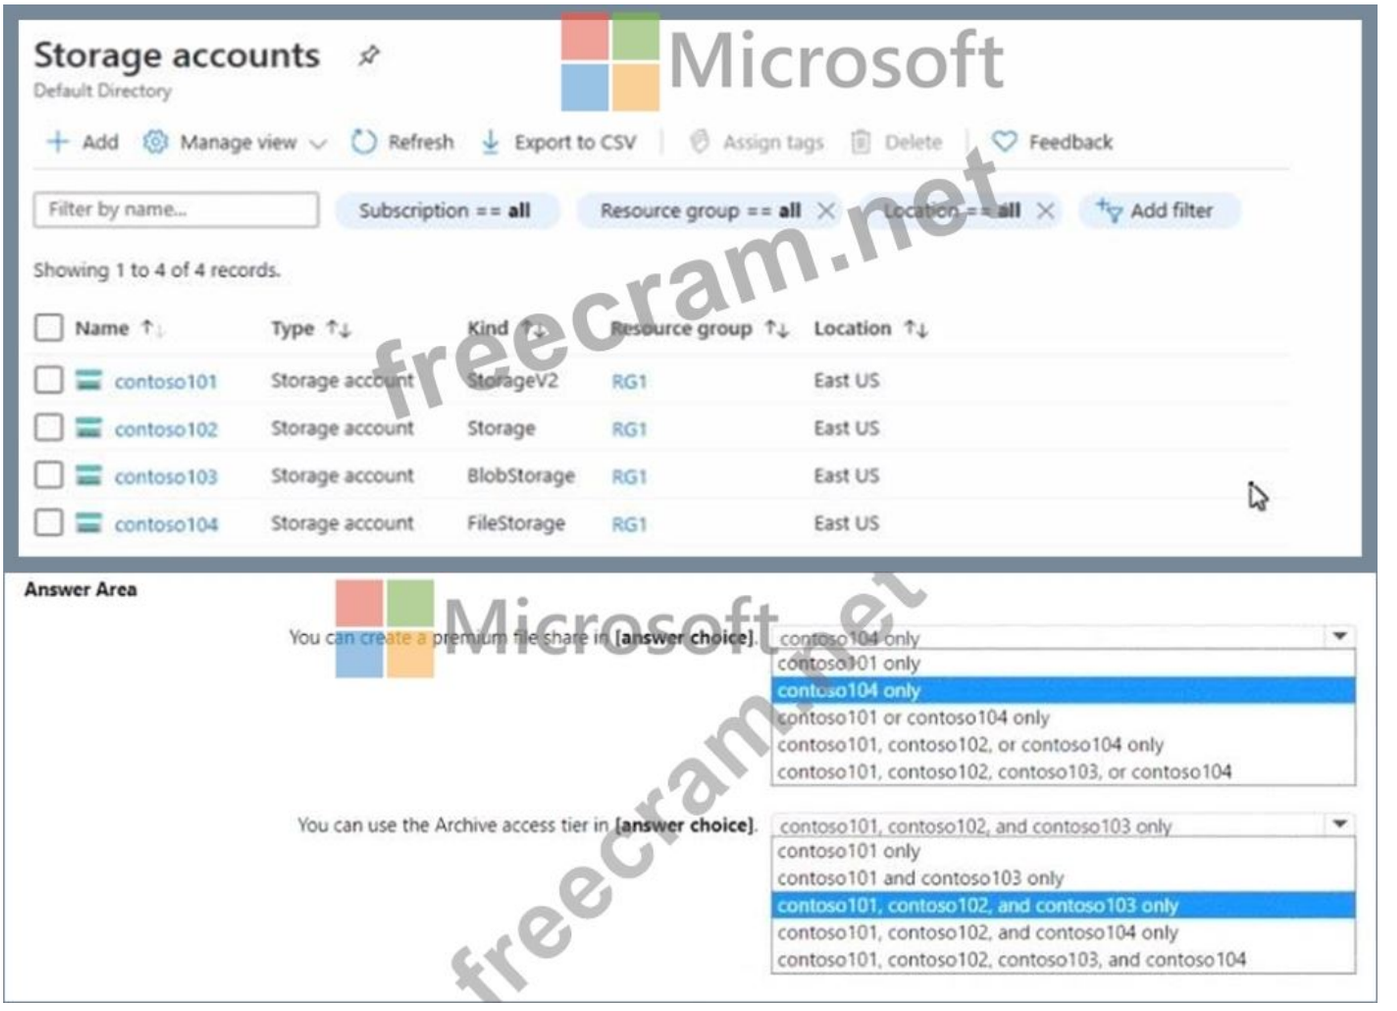

Answer:

Explanation:

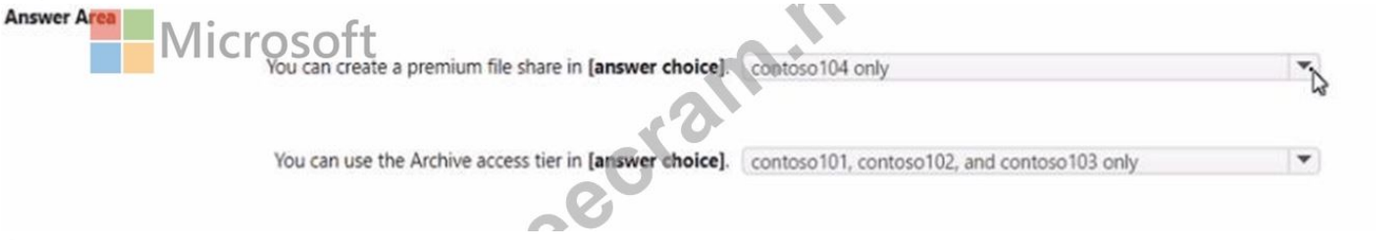

Scenario:

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: 136
You have an Azure subscription that contains the resources shown in the following table.

Name	Description
share1	File share in storage1
storage1	Storage account
User1	Microsoft Entra user

You need to assign User1 the Storage File Data SMB Share Contributor role for share1.
What should you do first?
A. Enable identity-based data access for the file shares instorage1.
B. Modify the security profile for the file shares in storage1.
C. Configure Access control (1AM) for share 1.
D. Select Default to Azure Active Directory authorization in the Azure portal for storage1.
Answer: (SHOW ANSWER)
When using an Azure Resource Manager (ARM) template to deploy multiple identical resources-such as several data disks for a virtual machine-you use the copy loop construct within the resource definition.
1. The Purpose of the copy Element
The copy element in an ARM template enables you to create multiple instances of a property or resource based on a defined count.
According to the Azure Resource Manager Template Schema Documentation:
"Use the copy element to repeat a resource property or resource definition multiple times during deployment.

The copy loop works with the copyIndex() function to generate a unique index value for each iteration." Therefore, the first selection should be copy, as it defines the structure that will be repeated for each data disk.

Example syntax:

```
"dataDisks": [
{
"copy": {
"name": "dataDisks",
"count": "[parameters('numberOfDataDisks')]",
"input": {
"lun": "[copyIndex()]",
"createOption": "Empty",
"diskSizeGB": 1023
}
}
}
]
```

2. The copyIndex() Function

The copyIndex() function returns the current iteration number within a copy loop (starting at 0 by default).

This allows each created disk to be assigned a unique Logical Unit Number (LUN) or a distinctive name.

Microsoft documentation states:

"The copyIndex() function returns the iteration index of a resource copy loop, which is often used to generate unique names or configuration values for each resource instance." Thus, the second selection (used to define lun) should be copyIndex(), ensuring each disk has a unique LUN value.

How It Works Together:

The copy block iterates based on the numberOfDataDisks parameter.

The copyIndex() function assigns each disk a unique identifier within the loop.

This structure ensures dynamic, scalable deployment of data disks without manually defining each one.

Final Verified Answer:

First Selection: copy

Second Selection: copyIndex()

Explanation Extracted from Microsoft Azure Administrator and ARM Template Documentation:

"The copy element repeats a property or resource in an ARM template."

"The copyIndex() function returns the index number of the iteration and can be used for unique naming or logical unit assignments." This combination (copy + copyIndex()) is the official and verified method for creating multiple data disks dynamically in an Azure virtual machine deployment using ARM templates.

NEW QUESTION: 137

You have an Azure subscription that contains the hierarchy shown in the following exhibit.

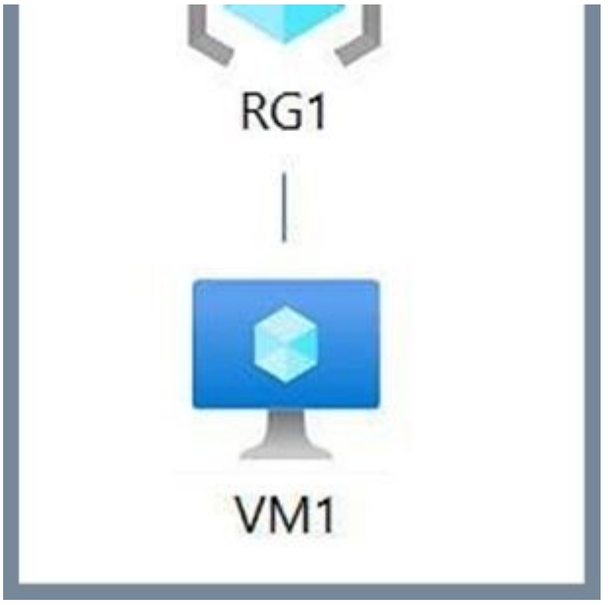

You create an Azure Policy definition named Policy1.

To which Azure resources can you assign Policy and which Azure resources can you specify as exclusions from Policy1? To answer, select the appropriate options in the answer NOTE Each correct selection is worth one point.

Answer Area

You can assign Policy1 to:

- Subscription1 and RG1 only
- ManagementGroup1 and Subscription1 only
- Tenant Root Group, ManagementGroup1, and Subscription1 only
- Tenant Root Group, ManagementGroup1, Subscription1, and RG1 only
- Tenant Root Group, ManagementGroup1, Subscription1, RG1, and VM1

You can exclude Policy1 from:

- VM1 only
- RG1 and VM1 only
- Subscription1, RG1, and VM1 only
- ManagementGroup1, Subscription1, RG1, and VM1 only
- Tenant Root Group, ManagementGroup1, Subscription1, RG1, and VM1

Answer:

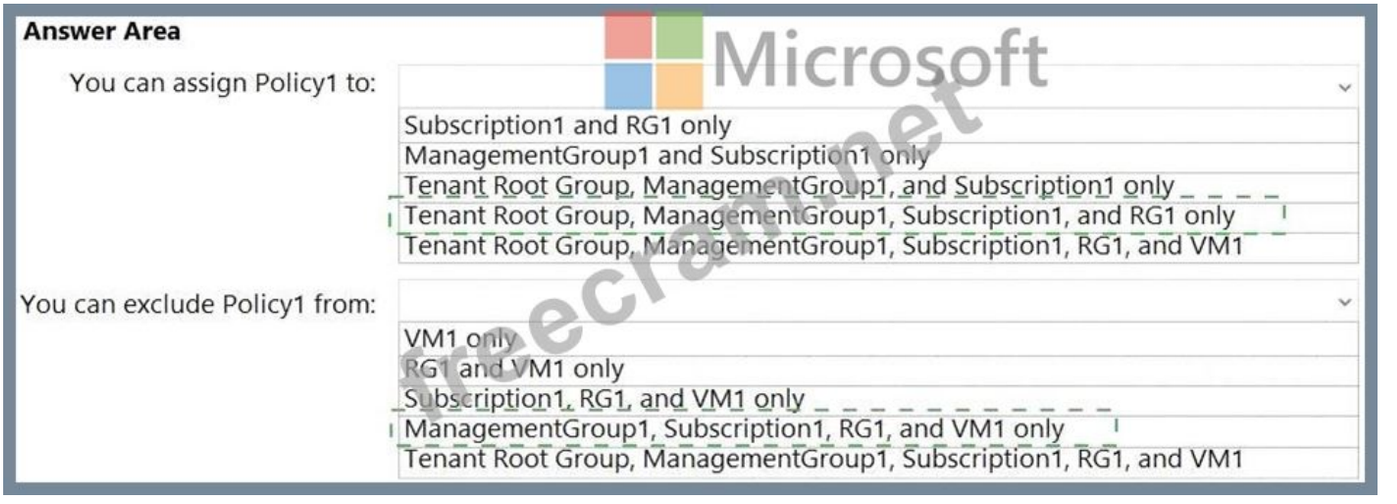

Explanation:

In Microsoft Azure, Azure Policy allows administrators to enforce rules and compliance standards

across management groups, subscriptions, resource groups, and individual resources. Policies can be assigned at multiple scopes, and the scope determines which resources the policy affects.
According to Microsoft Azure Governance and Policy documentation, an Azure Policy definition can be assigned to any of the following hierarchical scopes:
Tenant Root Group
Management Group
Subscription
Resource Group
All resources within the assigned scope, including nested resources (e.g., virtual machines within a resource group), inherit the policy. Therefore, a policy can be assigned at any level from the Tenant Root Group down to the Resource Group level, but not directly to individual resources such as a virtual machine (VM).
However, policy exclusions can be applied at any child scope under the assigned level. This means that if Policy1 is assigned at a higher level (for example, Tenant Root Group), you can exclude any lower-level scope, including:
A specific Subscription
A specific Resource Group
A specific Resource (e.g., VM)
Hence, Policy1 can be assigned at the Tenant Root Group, Management Group, Subscription, or Resource Group level, and excluded from lower levels (Subscription, RG1, or individual resources such as VM1).
Summary based on Microsoft Azure Policy design:
Assignable scopes: Tenant Root Group # Management Group # Subscription # Resource Group
Excludable scopes: Any child scope within the hierarchy, including specific resources (VMs).
Therefore:
You can assign Policy1 to: Tenant Root Group, ManagementGroup1, Subscription1, and RG1
You can exclude Policy1 from: Subscription1, RG1, and VM1 only

NEW QUESTION: 138
You have a virtual network named VNet1 that has the configuration shown in the following exhibit.

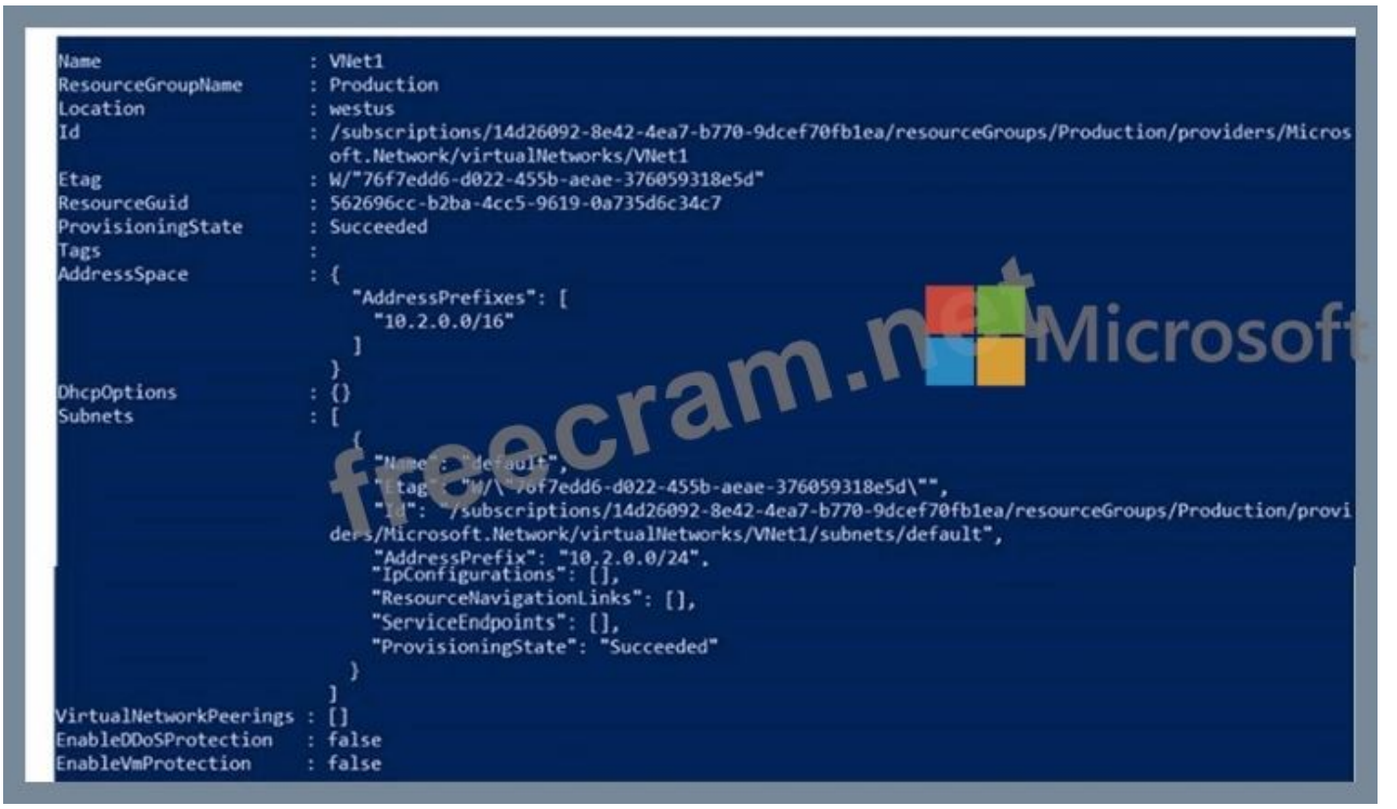

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

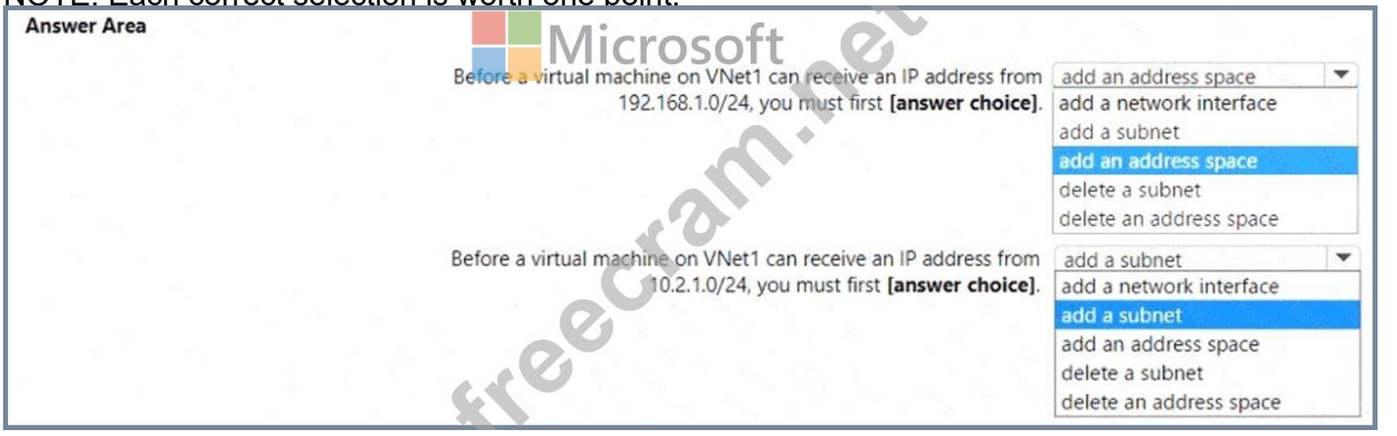

Answer:

Explanation:

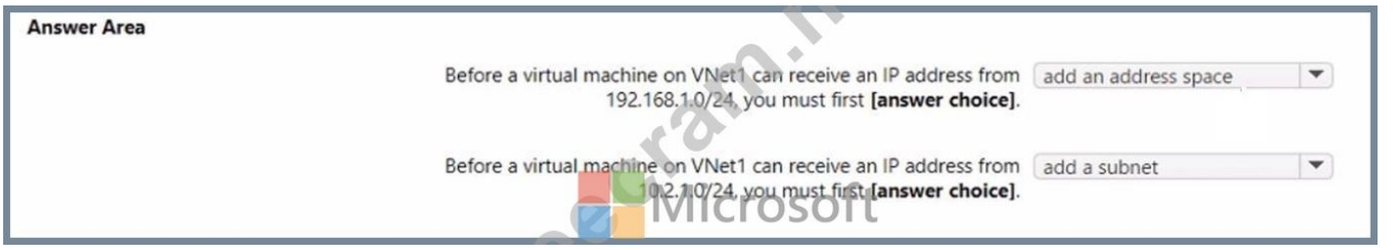

In this scenario, the PowerShell output shows the configuration of VNet1 in the West US region under the resource group Production. The AddressSpace for the virtual network is defined as 10.2.0.0/16, with a single existing subnet named default that has an AddressPrefix of 10.2.0.0/24.
According to Microsoft Azure networking documentation, an Azure virtual network (VNet) defines an address space (CIDR block), which represents the range of IP addresses available within that network.
Subnets are created within that address space to allocate smaller IP ranges for specific resource groupings, such as VMs or application tiers.
* For a VM to receive an IP address from 192.168.1.0/24, that range does not exist within the current VNet's address space (10.2.0.0/16). Therefore, before a VM can use that IP range, you must add an additional address space (192.168.1.0/24) to VNet1.
* For a VM to receive an IP address from 10.2.1.0/24, that range falls within the existing address space (10.2.0.0/16). However, there is no subnet currently defined for 10.2.1.0/24, so you must add a new subnet with that address prefix.
This configuration follows the official Azure documentation guidance:
"Each subnet must represent a subset of the virtual network's address space. To use a different IP range, you must first expand the VNet's address space before defining subnets within it."
Final Verified Answer:
* 192.168.1.0/24 # Add an address space
* 10.2.1.0/24 # Add a subnet

NEW QUESTION: 139
You need to ensure that you can grant Group4 Azure RBAC read-only permissions to all the A2ure file shares. What should you do?
A. On storagel and storage4, change the Account kind type to StorageV2 (general purpose v2).
B. Recreate storage2 and set Hierarchical namespace to Enabled.
C. On storage2, enable identity-based access for the file shares.
D. Create a shared access signature (SAS) for storagel, storage2, and storage4.
Answer: A (LEAVE A REPLY)
Azure role-based access control (Azure RBAC) for Azure file shares requires identity-based authentication integration. According to the Microsoft Azure Administrator documentation, this feature is only supported for StorageV2 (general purpose v2) and FileStorage account types.
In this scenario:
You are required to grant Group4 read-only access using Azure RBAC on Azure file shares.
The technical requirement specifies:
"Whenever possible, grant Group4 Azure RBAC read-only permissions to the Azure file shares."

From the case study data:
Storage Account
Kind
Identity-based Access
storage1
Storage (general purpose v1)
Azure AD DS
storage2
StorageV2
Disabled
storage3
BlobStorage
N/A
storage4
FileStorage
Azure AD DS
The Storage (general purpose v1) type (storage1) does not support Azure AD or Azure RBAC integration for file shares. Microsoft documentation clearly states that "StorageV1 accounts must be upgraded to StorageV2 to support Azure AD authentication and RBAC role assignments."
Meanwhile, FileStorage (storage4) already supports Azure AD Domain Services (Azure AD DS) and RBAC role assignment; hence no further modification is required there. However, to make storage1 compatible, it must be converted from StorageV1 to StorageV2.
Once converted to StorageV2, you can then:
Enable identity-based access for Azure file shares.
Assign Azure RBAC roles (e.g., Storage File Data Reader) to Group4.
Microsoft-Documented Requirements Summary:
Supported Account Types: StorageV2 or FileStorage
Unsupported: StorageV1 and BlobStorage
Required RBAC Roles for Read-Only Access:
Storage File Data Reader (or custom read-only role)
Thus, to meet the organization's requirement to provide Azure RBAC read-only permissions, you must change the account type of storage1 to StorageV2, ensuring both storage1 and storage4 can be managed with Azure RBAC.
Topic 4, Humongous InsuranceOverview
Existing Environment
Huongous Insurance is an insurance company that has three offices in Miami, Tokoyo, and Bankok. Each has
5000 users.
Active Directory Environment
Humongous Insurance has a single-domain Active Directory forest named humongousinsurance.com. The functional level of the forest is Windows Server 2012.

You recently provisioned an Azure Active Directory (Azure AD) tenant.
Network Infrastructure
Each office has a local data center that contains all the servers for that office. Each office has a dedicated connection to the Internet.
Each office has several link load balancers that provide access to the servers.
Active Directory Issue
Several users in humongousinsurance.com have UPNs that contain special characters.
You suspect that some of the characters are unsupported in Azure AD.
Licensing Issue
You attempt to assign a license in Azure to several users and receive the following error message: "Licenses not assigned. License agreement failed for one user." You verify that the Azure subscription has the available licenses.
Requirements
Planned Changes
Humongous Insurance plans to open a new office in Paris. The Paris office will contain 1,000 users who will be hired during the next 12 months. All the resources used by the Paris office users will be hosted in Azure.
Planned Azure AD Infrastructure
The on-premises Active Directory domain will be synchronized to Azure AD.
All client computers in the Paris office will be joined to an Azure AD domain.
Planned Azure Networking Infrastructure
You plan to create the following networking resources in a resource group named All_Resources: Default Azure system routes that will be the only routes used to route traffic A virtual network named Paris-VNet that will contain two subnets named Subnet1 and Subnet2 A virtual network named ClientResources-VNet that will contain one subnet named ClientSubnet A virtual network named AllOffices-VNet that will contain two subnets named Subnet3 and Subnet4 You plan to enable peering between Paris-VNet and AllOffices-VNet. You will enable the Use remote gateways setting for the Paris-VNet peerings.
You plan to create a private DNS zone named humongousinsurance.local and set the registration network to the ClientResources-VNet virtual network.
Planned Azure Computer Infrastructure
Each subnet will contain several virtual machines that will run either Windows Server 2012 R2, Windows Server 2016, or Red Hat Linux.
Department Requirements
Humongous Insurance identifies the following requirements for the company's departments:
Web administrators will deploy Azure web apps for the marketing department. Each web app will be added to a separate resource group. The initial configuration of the web apps will be identical.
The web administrators have permission to deploy web apps to resource groups.
During the testing phase, auditors in the finance department must be able to review all Azure costs from the past week.
Authentication Requirements

Users in the Miami office must use Azure Active Directory Seamless Single Sign-on (Azure AD Seamless SSO) when accessing resources in Azure.

NEW QUESTION: 140

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You need to ensure that an Azure Active Directory (Azure AD) user named Admin1 is assigned the required role to enable Traffic Analytics for an Azure subscription.
Solution: You assign the Network Contributor role at the subscription level to Admin1.
Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

Your account must meet one of the following to enable traffic analytics:
Your account must have any one of the following Azure roles at the subscription scope: owner, contributor, reader, or network contributor.
Reference:
https://docs.microsoft.com/en-us/azure/network-watcher/traffic-analytics-faq

NEW QUESTION: 141

You have an Azure virtual machine named VM1 and an Azure key vault named Vault1.
On VM1, you plan to configure Azure Disk Encryption to use a key encryption key (KEK) You need to prepare Vault! for Azure Disk Encryption.
Which two actions should you perform on Vault1? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Create a new key.

B. Select Azure Virtual machines for deployment

C. Configure a key rotation policy.

D. Create a new secret.

E. Select Azure Disk Encryption for volume encryption

Answer: (SHOW ANSWER)

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: 142

You have an Azure subscription that contains the resources shown in the following table.

Name	Location	Description
RG1	East US	Resource group
ASP1	East US	Azure App Service plan
ASP2	West Europe	Azure App Service plan
WebApp1	East US	Azure App Service web app in ASP1
Managed1	East US	Managed identity

WebApp1 has the following configurations:

* Identity: Managed!
* Environment variable: EnvVar1
* Deployment slots: WebApp1 -slot1
* Backups: Perform a custom backup every two days

You clone WebApp1 by running the following commands.

```
$srcapp = Get-AzWebApp -ResourceGroupName RG1 -Name WebApp1
New-AzWebApp -ResourceGroupName RG1 -Name WebApp2 -Location "West Europe" -AppServicePlan ASP2 -SourceWebApp $srcapp
```

Which configuration of WebApp1 is cloned to WebApp2?

A. Identity

B. Deployment slots

C. Backups

D. Environment variable

Answer: (SHOW ANSWER)

When you clone an Azure App Service web app using New-AzWebApp -SourceWebApp, Azure copies application settings and configuration that are part of the app's configuration metadata. Environment variables (App Settings) are included in the clone.

However, the following are not cloned:

* Managed identity (must be enabled separately on the new app)
* Deployment slots (slots are not copied to the new web app)
* Backup configuration (backup schedules and settings are not cloned)

Therefore, among the listed options, only the environment variable is cloned to WebApp2.

NEW QUESTION: 143

You need to add VM1 and VM2 to the backend poo! of LB1. What should you do first?

A. Create a new NSG and associate the NSG to VNET1/Subnet1.

B. Connect VM2 to VNET1/Subnet1.

C. Redeploy VM1 and VM2 to the same availability zone.

D. Redeploy VM1 and VM2 to the same availability set.

Answer: (SHOW ANSWER)

In Azure, Load Balancers distribute network traffic across multiple virtual machines (VMs) to ensure high availability and reliability. To add virtual machines to the backend pool of an Azure

Load Balancer, the following key conditions must be met according to the Microsoft Azure Administrator documentation:
All VMs in the backend pool must be connected to the same virtual network (VNet) as the Load Balancer.
The Load Balancer (in this case, LB1) is configured for internal load balancing on VNET1/Subnet1 as per the technical requirements of the case study.
The backend pool can include network interfaces (NICs) from VMs within the same region and VNet.
Step-by-step analysis:
From the case study data:
VM
Location
Connected to
IP Address
VM1
West US
VNET1/Subnet1
10.0.1.4
VM2
West US
VNET1/Subnet2
10.0.2.4
LB1
Internal Basic Load Balancer
Connected to VNET1/Subnet1
-
Observation:
VM1 is already connected to VNET1/Subnet1, where the internal Load Balancer LB1 is also deployed.
VM2, however, is connected to VNET1/Subnet2, which is a different subnet within the same virtual network.
According to Microsoft Learn ("Configure backend pools in Azure Load Balancer"):
"All network interfaces in the backend pool must be within the same virtual network as the load balancer.
You cannot add VMs connected to different VNets or subnets not associated with the load balancer's front- end configuration." Therefore, before you can add VM2 to the backend pool, you must ensure that its network interface is attached to VNET1/Subnet1, the same subnet used by LB1.
Only after this step will both VMs (VM1 and VM2) be eligible for inclusion in LB1's backend pool.
Incorrect Option Analysis:
A). Create a new NSG and associate the NSG to VNET1/Subnet1.

Not required. Network Security Groups control traffic filtering, not backend pool configuration.
C). Redeploy VM1 and VM2 to the same availability zone.
Availability Zones only matter for redundancy and failover, not for backend pool eligibility in a basic internal load balancer.
D). Redeploy VM1 and VM2 to the same availability set.
Basic Load Balancers can distribute traffic across VMs in the same availability set, but both VMs must already reside in the same VNet/Subnet first.
Final Verified Answer:
B. Connect VM2 to VNET1/Subnet1
Reference (Microsoft Official Documentation):
Microsoft Learn: Configure the backend pool for Azure Load Balancer
Microsoft Learn: Azure Load Balancer overview
Microsoft Learn: Create and configure an internal load balancer
Microsoft Learn: Virtual network and subnet requirements for load balancing

NEW QUESTION: 144
You have an Azure Storage account named storage1.
You need to enable a user named User1 to list and regenerate storage account keys for storage1.
Solution: You assign the Storage Account Contributor role to User1.
Does this meet the goal?
A. Yes
B. No
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: 145
You have an Azure subscription
You plan to deploy a new storage account
You need to configure encryption for the account The solution must meet the following requirements

* Use a customer-managed key stored in an key vault
* Use the maximum supported bit length.
Which type of key and which bit length should you use?

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: 146
You have an Azure subscription that contains a resource group named RG26.
RG26 is sot to the West Europe location and is used to create temporary resources for a project.
RG26 contains the resources shown in the following table.

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 is backed up to RGV1.

When the project is complete, you attempt to delete RG26 from the Azure portal. The deletion fails.

You need to delete RG26.

What should you do first?

A. Stop the backup of SQLDB01.

B. Delete sa001.

C. Delete VM1.

D. StopVM1.

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: 147

You need to define a custom domain name for Azure AD to support the planned infrastructure.

Which domain name should you use?

A. Join the client computers in the Miami office to Azure AD.

B. Add http://autologon.microsoftazuread-sso.com to the intranet zone of each client computer in the Miami office.

C. Allow inbound TCP port 8080 to the domain controllers in the Miami office.

D. Install Azure AD Connect on a server in the Miami office and enable Pass-through Authentication

E. Install the Active Directory Federation Services (AD FS) role on a domain controller in the Miami office.

Answer: (SHOW ANSWER)

Every Azure AD directory comes with an initial domain name in the form of domainname.onmicrosoft.com.

The initial domain name cannot be changed or deleted, but you can add your corporate domain name to Azure AD as well. For example, your organization probably has other domain names used to do business and users who sign in using your corporate domain name. Adding custom domain names to Azure AD allows you to assign user names in the directory that are familiar to

your users, such as 'alice@contoso.com.' instead of
'alice@domain name.onmicrosoft.com'.
Scenario:
Network Infrastructure: Each office has a local data center that contains all the servers for that office. Each office has a dedicated connection to the Internet.
Humongous Insurance has a single-domain Active Directory forest named humongousinsurance.com Planned Azure AD Infrastructure: The on-premises Active Directory domain will be synchronized to Azure AD.
References: https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/add-custom-domain

NEW QUESTION: 148
You have an Azure virtual machine named VM1. VM1 was deployed by using a custom Azure Resource Manager template named ARMIjson.
You receive a notification that VM1 will be affected by maintenance.
You need to move VM1 to a different host immediately.
Solution: From the VM1 Updates blade, select One-time update.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
When a virtual machine (VM) in Azure is scheduled for platform maintenance, it can either be planned maintenance (with notification) or unplanned maintenance (without notice). During planned maintenance, administrators have the option to move the VM to a different host proactively to minimize downtime.
According to Microsoft Azure documentation ("Maintenance and updates for Azure virtual machines"), the correct action is to use the "Redeploy" or "Move to another host" feature, not a "One-time update." The
"Redeploy" operation deallocates the VM and migrates it to a new node within Azure's infrastructure, resolving underlying host issues while preserving OS and data disks.
The "Updates blade" in the Azure portal does not contain a "One-time update" option to move a VM off a host. That term is used in the context of patching or guest OS updates, not host maintenance.
Therefore, the proposed solution does not meet the goal of moving the VM to another host immediately. The correct action would be to redeploy the VM using the Azure portal or Invoke-AzVMRedeploy PowerShell cmdlet.

NEW QUESTION: 149
You have the Azure virtual machines shown in the following table.

Name	IP address	Virtual network
VM1	10.0.0.4	VNET1
VM2	10.0.0.5	VNET1

VNET1 is linked to a private DNS zone and named contoso.com that contains the records shown in the following table.

Name	Type	TTL	Value	Auto registered
comp1	TXT	3600	10.0.0.5	False
comp2	A	3600	10.0.0.5	False
comp3	CNAME	3600	comp1.contoso.com	False
comp4	PTR	3600	10.0.0.5	False

You need to ping VM2 from VM1.

Which DNS names can you use to ping VM2.

A. comp1.contoso.com comp2contoso.com.comp3.contoso.com and comp4.contoso.com

B. comp2.contoso.com and comp4.contoso.com only

C. comp1.contoso.com, comp2.contoso.com and comp4.contoso.con only

D. comp2 contoso.com only

E. com1.contoso.com and comp2.contoso.com only

Answer: (SHOW ANSWER)

NEW QUESTION: 150

You have an Azure subscription that contains the resources shown in the following table.

Name	Type
LB1	Load balancer
VM1	Virtual machine
VM2	Virtual machine

LB1 is configured as shown in the following table.

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

You plan to create new inbound NAT rules that meet the following requirements:

Provide Remote Desktop access to VM2 from the internet by using port 3389.

A. A frontend IP address

B. A health probe

C. A load balancing rule

D. A backend pool

Answer: (SHOW ANSWER)

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: 151
You have an Azure subscription.
You need to receive an email alert when a resource lock is removed from any resource in the subscription What should you use to create an activity log alert in Azure Monitor?
A. a resource, a condition, and an action group
B. a resource, a condition, and a Microsoft 365 group
C. a Log Analytics workspace, a resource, and an action group
D. a data collection endpoint, an application security group, and a resource group
Answer: (SHOW ANSWER)
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: 152
You plan to deploy several Azure virtual machines that will run Windows Server 2022 in a virtual machine scale set by using an Azure Resource Manager template.
You need to ensure that NGINX is available on all the virtual machines after they are deployed.
What should you use?
A. Azure Application Insights
B. Azure Custom Script Extension
C. the Publish-ArVMDscConfiguration cmdlet
D. the New-AzConfigurationAssignment Cmdlet
Answer: (SHOW ANSWER)
When deploying Azure Virtual Machine Scale Sets (VMSS) using an Azure Resource Manager (ARM) template, you can automatically install and configure applications or services (like NGINX) on the virtual machines using the Azure Custom Script Extension.
According to the Microsoft Azure Administrator documentation, the Custom Script Extension is used to download and execute scripts on Azure virtual machines after deployment. It is ideal for configuration management, application installation, and post-deployment software setup.
The extension can execute PowerShell or Bash scripts stored in Azure Storage or GitHub repositories. When used with an ARM template, you define the extension in the template's "resources" section under "type": "Microsoft.Compute/virtualMachines/extensions" or "Microsoft.Compute/virtualMachineScaleSets/extensions".
Microsoft's official documentation notes:
"The Custom Script Extension for Windows and Linux downloads and executes scripts on Azure virtual machines. This extension is useful for post-deployment configuration, software installation, or any other configuration or management tasks." In this case, using Azure Custom Script Extension allows automatic installation of NGINX on all VM instances as soon as they are created in the scale set - satisfying the requirement that NGINX is available immediately after deployment.
Other options do not fit:
* Application Insights is used for monitoring, not configuration.
* Publish-AzVMDscConfiguration and New-AzConfigurationAssignment are used for Desired State Configuration (DSC), not lightweight application installation like NGINX.
Hence, the only verified and cost-efficient solution is Azure Custom Script Extension.

NEW QUESTION: 153

You have an Azure Subscription that contains the virtual networks Shown in the following table.

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

All the virtual networks are peered. Each virtual network contains nine virtual machines.
You need to configure secure RDP corrections to the virtual machines by using Azure Boston.
Whit is the minimum number of Bastion nests required?

A. 1

B. 3

C. 9

D. 10

Answer: (SHOW ANSWER)

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: 154

You have an Azure subscription.
You plan to deploy the Azure container instances shown in the following table.

Name	Operating system
Instance1	Nano Server installation of Windows Server 2019
Instance2	Server Core installation of Windows Server 2019
Instance3	Linux
Instance4	Linux

Which instances can you deploy to a container group?

A. Instance1 only

B. Instance2only

C. Instance1 and Instance2 only

D. Instance3 and Instance4 only

Answer: (SHOW ANSWER)

https://learn.microsoft.com/en-us/azure/container-instances/container-instances-container-groups Multi- container groups currently support only Linux containers. For Windows containers, Azure Container Instances only supports deployment of a single container instance. While we are working to bring all features to Windows containers, you can find current platform differences in the service

NEW QUESTION: 155

You have an Azure subscription that contains a virtual machine named VM1.

To VM1, you plan to add a 1-TB data disk that meets the following requirements:

* Provides data resiliency in the event of a datacenter outage.
* Provides the lowest latency and the highest performance.
* Ensures that no data loss occurs if a host fails.

You need to recommend which type of storage and host caching to configure for the new data disk.

Answer Area

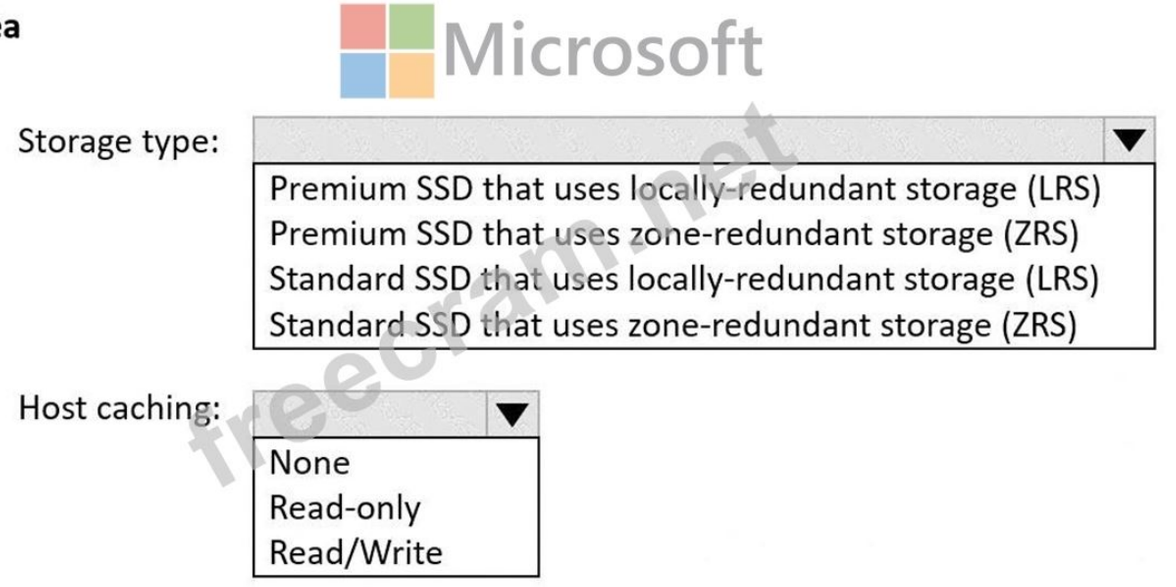

Answer:

Answer Area

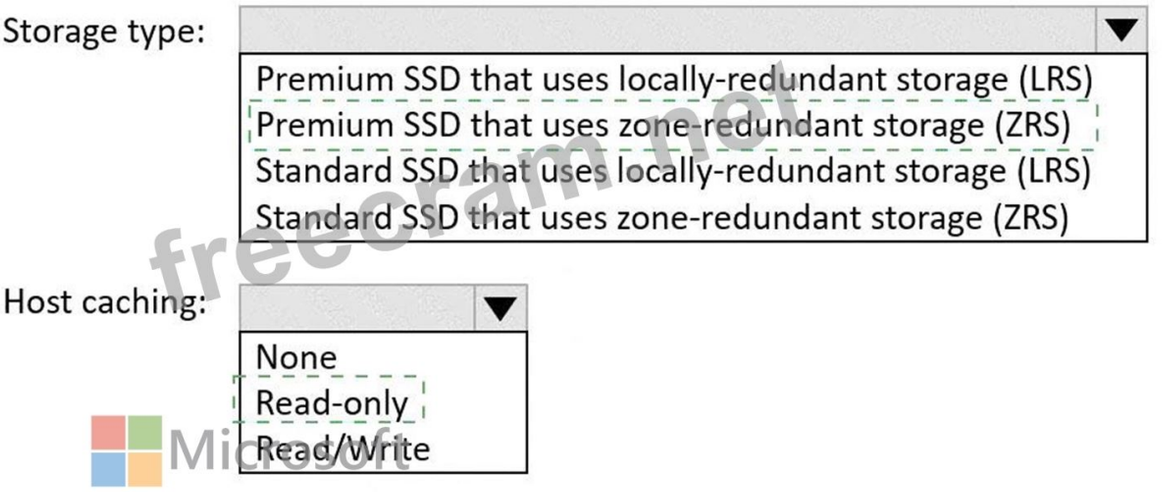

Explanation:

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: 156

You have an Azure subscription.

You deploy a virtual machine scale set that is configure as shown in the following exhibit.

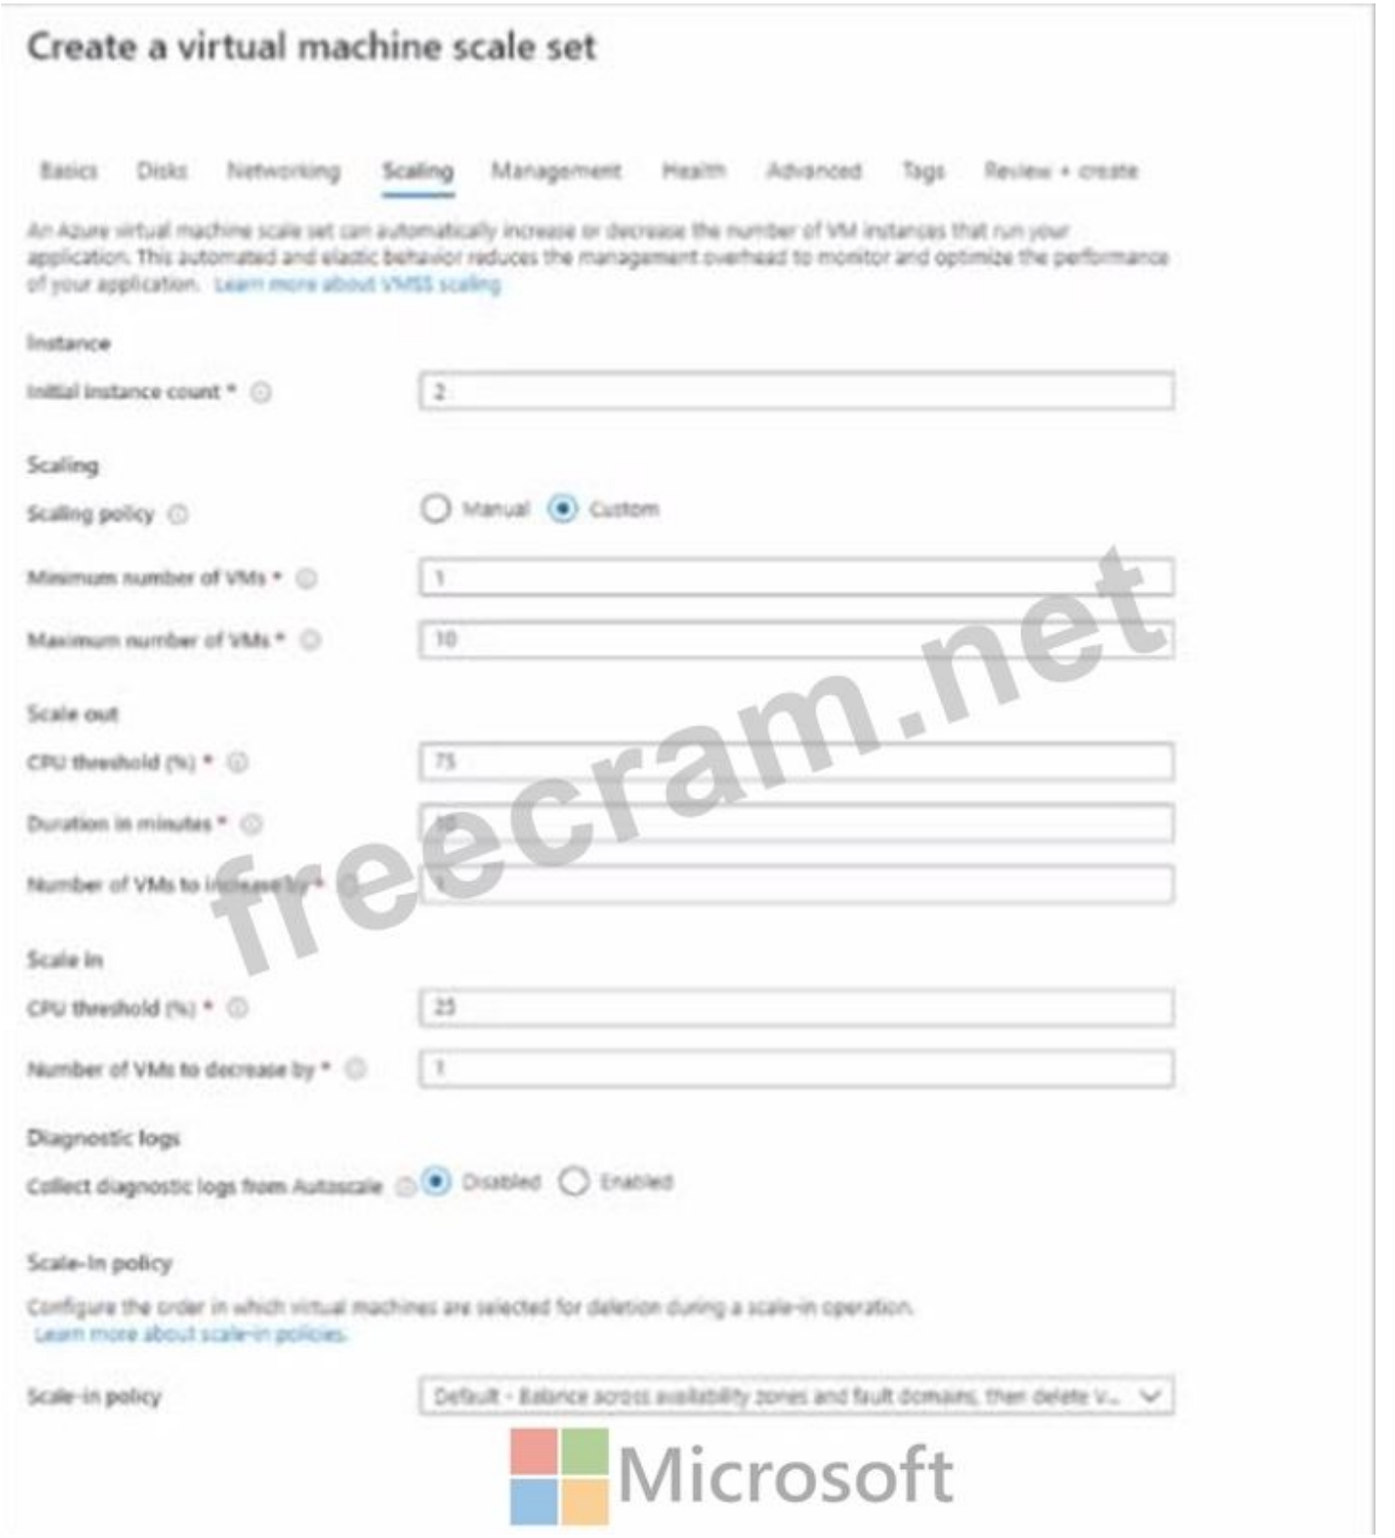

Use the drop-down menus to select the answer choice that answers each questions based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

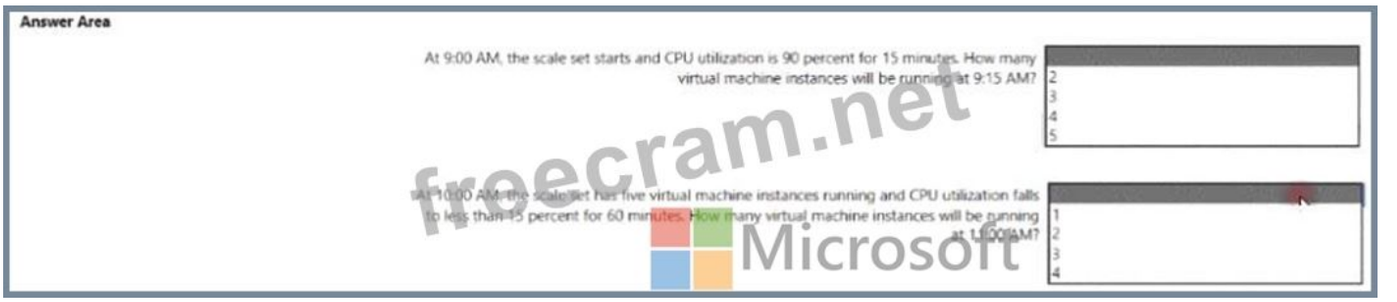

Answer:

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: 157

You have an Azure subscription that contains a storage account named storage1.

You need to configure a shared access signature (SAS) to ensure that users can only download blobs securely by name.

Which two settings should you configure? To answer, select the appropriate settings in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

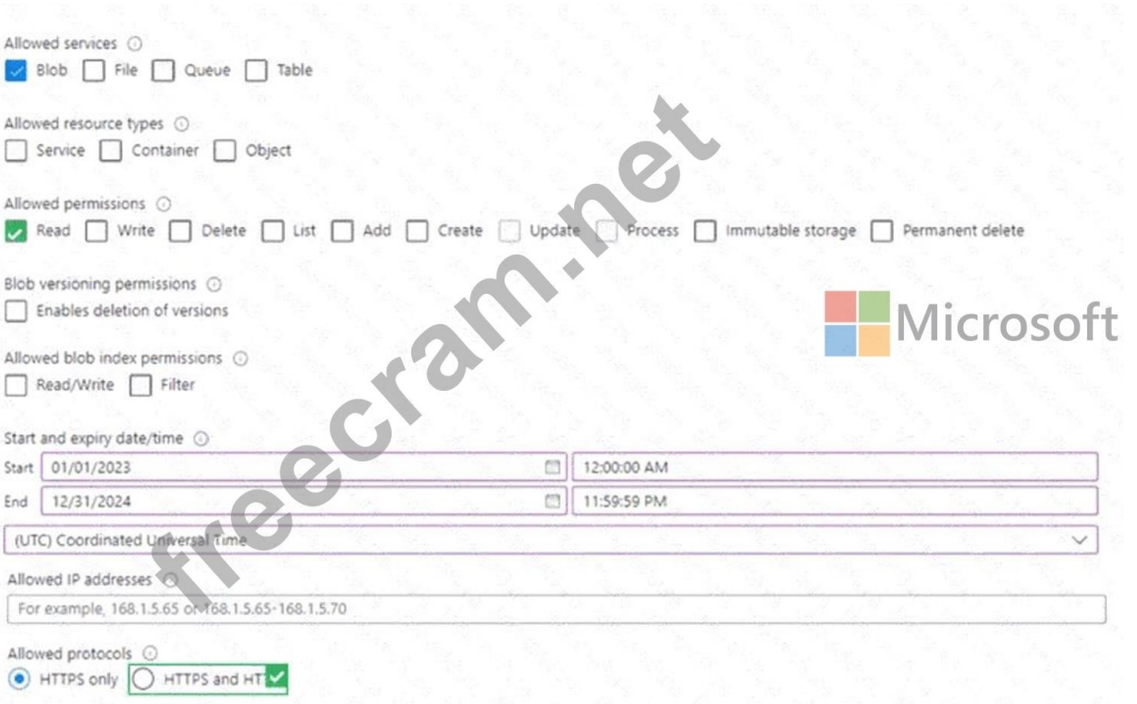

Answer:

Answer Area

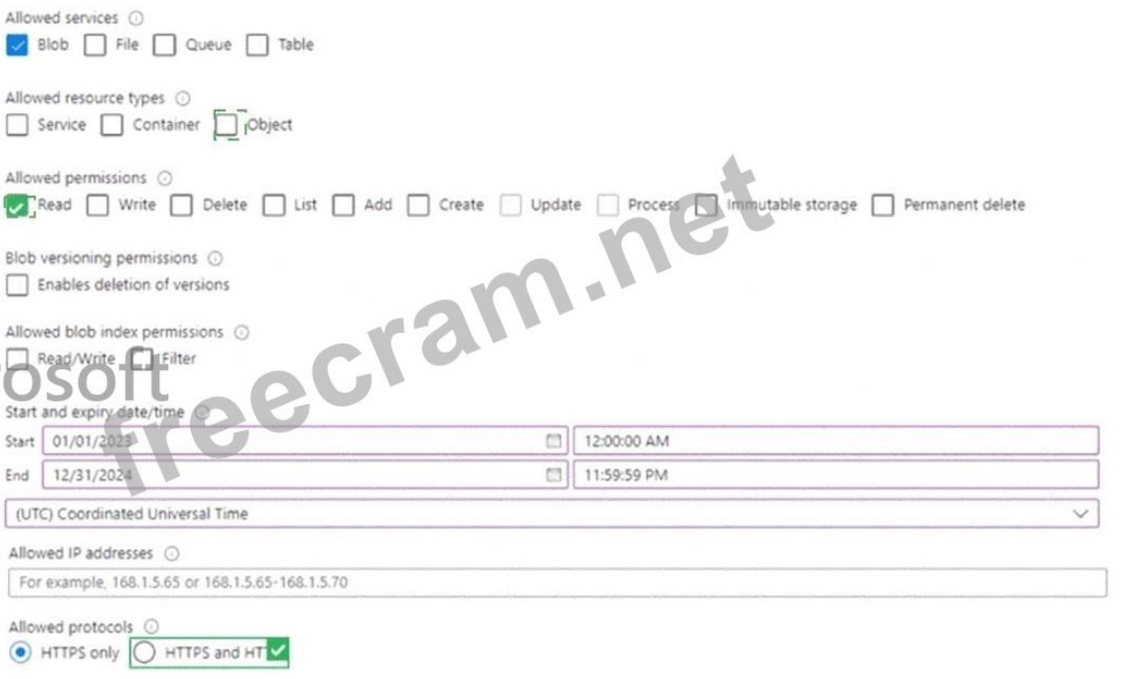

NEW QUESTION: 158

You need to ensure that VM1 can communicate with VM4. The solution must minimize administrative effort.

What should you do?

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

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

C. Establish peering between VNET1 and VNET3.

D. Create an NSG and associate the NSG to VMI and 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: 159
You need to configure the alerts for VM1 and VM2 to meet the technical requirements.
Which three actions should you perform in sequence? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

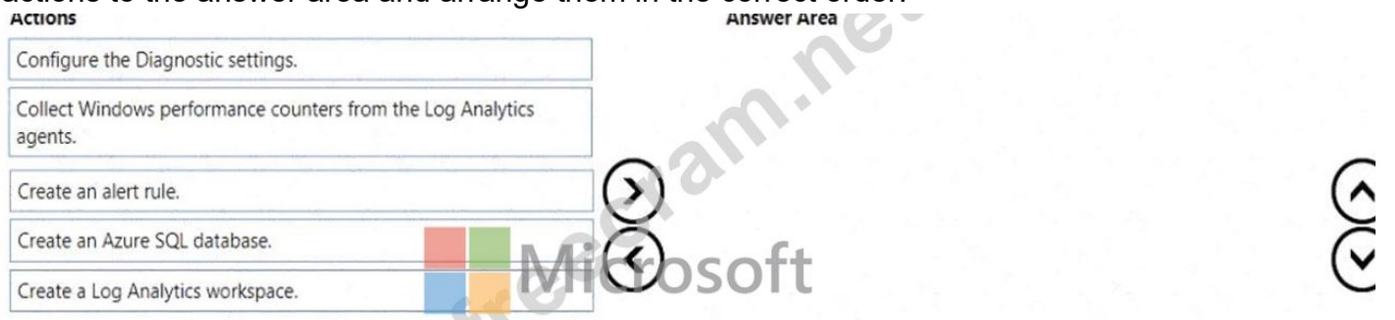

Answer:

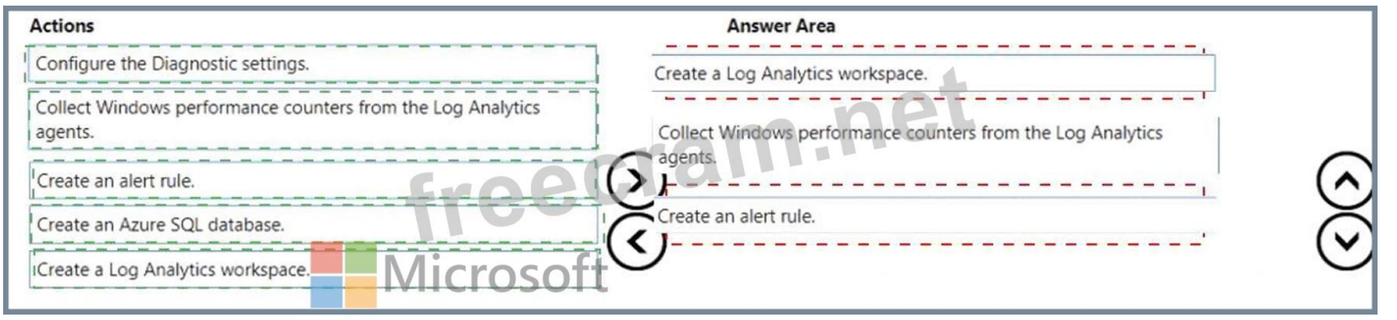

Explanation:
1## Create a Log Analytics workspace.
2## Collect Windows performance counters from the Log Analytics agents.
3## Create an alert rule.
The question states:
"You need to configure the alerts for VM1 and VM2 to meet the technical requirements." The technical requirement from the case study 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: 160
You need to identify the storage requirements for Contoso.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

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: 161

You have an Azure subscription named Subscription1 that contains the virtual networks in the following table.

Name	Subnet
VNet1	Sybnet11
VNet2	Subnet12
VNet3	Subnet13

Subscripton1 contains the virtual machines in the following table.

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

In Subscription1, you create a load balancer that has the following configurations:

Name: LB1

SKU: Basic

Type: Internal

Subnet: Subnet12

Virtual network: VNET1

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: each correct selection is worth one point.

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: 162
You have an Azure subscription that has a Recovery Services vault named Vault 1. The subscription contains the virtual machines shown in the following table.

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

You plan to schedule backups to occur every night at 23:00.

Which virtual machines can you back up by using Azure Backup?

A. VM1, VM2, VM3 and VM4

B. VM1 and VM3 only

C. VM1 only

D. VM1 and VM2 only

Answer: (SHOW ANSWER)

NEW QUESTION: 163

You have the role assignment file shown in the following exhibit.

```
[
    {
        "RoleAssignmentId": "e3108585-0e5d-4572-91a3-aa5d2df73999",
        "Scope": "/subscriptions/fb960108-fcdc-499b-886e-d9c31d3f26ff",
        "DisplayName": "User1",
        "SignInName": "User1@contoso.onmicrosoft.com",
        "RoleDefinitionName": "Owner",
        ...
    },
    {
        "RoleAssignmentId": "3bab4763-16a9-4d5d-9fcd-eee0cc31a21e",
        "Scope": "/subscriptions/fb960108-fcdc-499b-886e-d9c31d3f26ff/resourceGroups/RG2",
        "DisplayName": "User2",
        "SignInName": "User2@contoso.onmicrosoft.com",
        "RoleDefinitionName": "Owner",
        ...
    },
    {
        "RoleAssignmentId": "a071c023-40a3-4b7f-8680-1109b40270c5",
        "Scope": "/subscriptions/fb960108-fcdc-499b-886e-d9c31d3f26ff/resourceGroups/RG1/providers/
Microsoft.Compute/virtualMachines/VM1",
        "DisplayName": "User3",
        "SignInName": "User3@contoso.onmicrosoft.com",
        "RoleDefinitionName": "Owner",
        ...
    },
    {
        "RoleAssignmentId": "c5b9e7da-76d4-4888-93b5-8afb2bb780b4",
        "Scope": "/subscriptions/fb960108-fcdc-499b-886e-d9c31d3f26ff/resourceGroups/RG1",
        "DisplayName": "User4",
        "SignInName": "User4@contoso.onmicrosoft.com",
        "RoleDefinitionName": "Contributor",
        ...
    }
]
```


Use the drop-down menus to select the answer choice that completes

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worm one point.

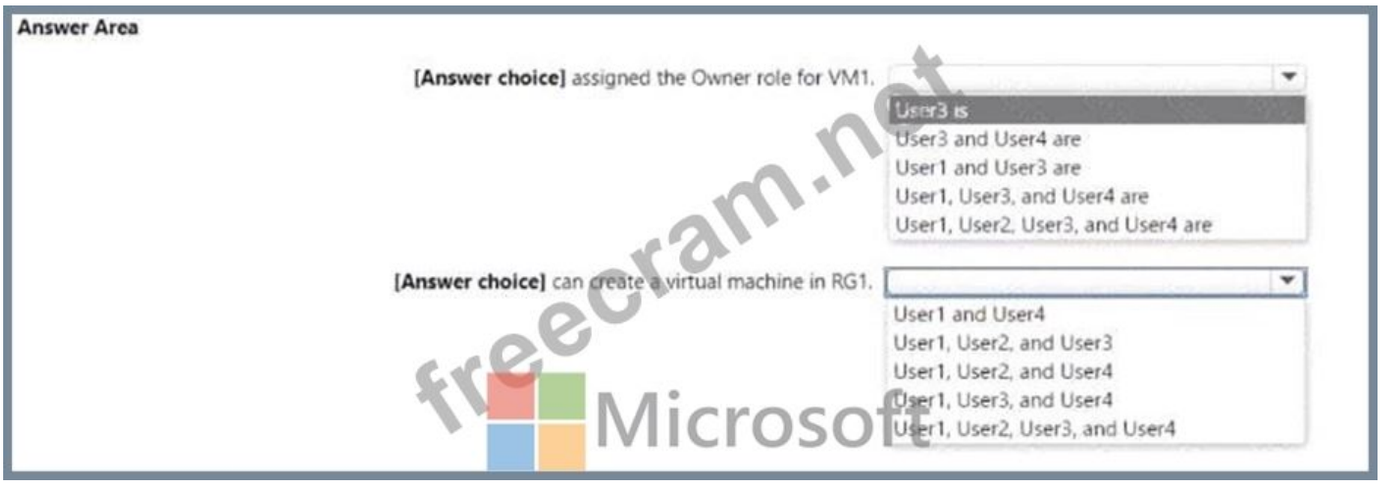

Answer:

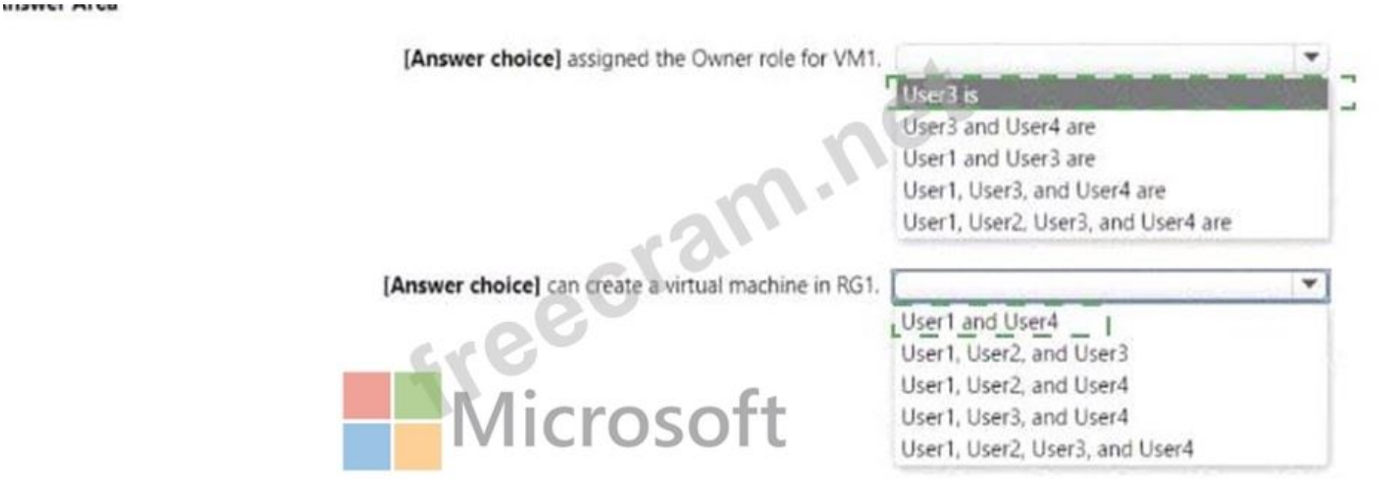

Explanation:

User3 is assigned the Owner role for VM1.

User1, and User4 can create a virtual machine in RG1.

NEW QUESTION: 164

You have an Azure subscription named Sub1 that contains the blob containers shown in the following table.

Sub1 contains two users named User1 and User2. Both users are assigned the Reader role at the Sub1 scope.

You have a condition named Condition1 as shown in the following exhibit.

```
(
 (
  !(ActionMatches{'Microsoft.Storage/storageAccounts/blobServices/containers/blobs/read'})
 )
 OR
 (
  @Resource[Microsoft.Storage/storageAccounts/blobServices/containers:name] StringEquals 'cont1'
 )
)
```

You have a condition named Condition2 as shown in the following exhibit.

```
(
 (
  !(ActionMatches{'Microsoft.Storage/storageAccounts/blobServices/containers/blobs/write'})
 )
 OR
 (
  @Resource[Microsoft.Storage/storageAccounts/blobServices/containers/blobs:path] StringLike '*2*'

 )
)
```

You assign roles to User1 and User2 as shown in the following table.

User	Role	Scope	Role assignment condition
User1	Storage Blob Data Reader	Sub1	Condition1
User2	Storage Blob Data Owner	storage1	Condition2

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

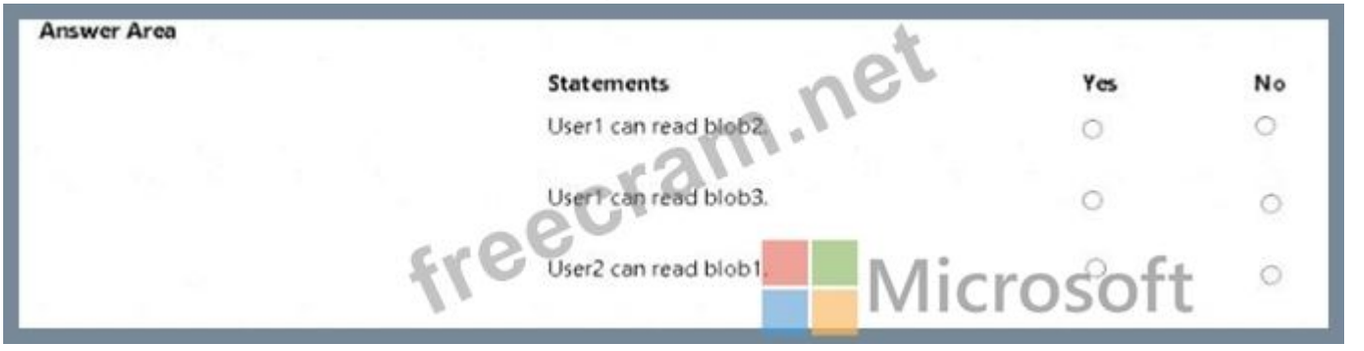

Answer:

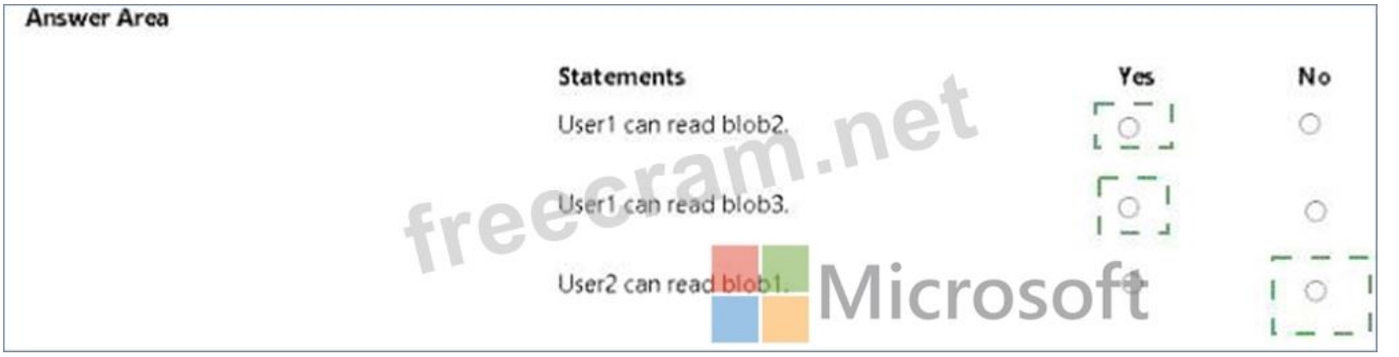

Explanation:

Answer Area

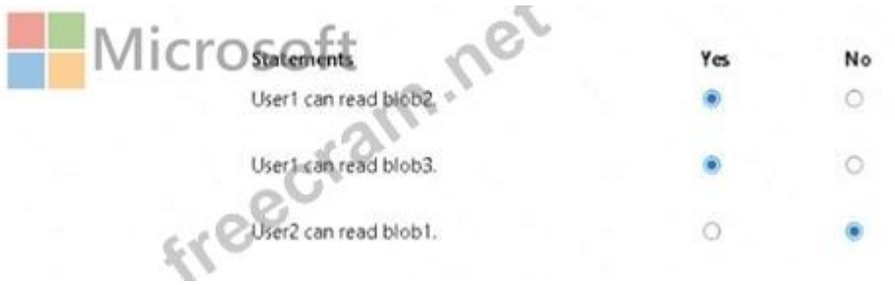

NEW QUESTION: 165

You have an Azure subscription that contains the virtual networks shown in the following table.

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

You need to deploy an Azure firewall named AF1 to RG1 in the West US Azure region.

To which virtual networks can you deploy AF1?

A. VNET1 only

B. VNET1 and VNET2 only

C. VNET1 and VNET4 only

D. VNET1, VNET2, and VNET4 only

E. VNET1, VNET2, VNET3, and VNET4

Answer: (SHOW ANSWER)

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: 166

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

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

You have an Azure Storage account named storage1.

You need to enable a user named User1 to list and regenerate storage account keys for storage I.

Solution: You assign the Reader and Data Access role to User1.

Does this meet the goal?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

Azure Storage account access keys are highly privileged credentials that provide full control over the storage account, including access to all data and the ability to regenerate keys. Because of the security implications, listing and regenerating storage account keys is classified as a management-plane operation, not a data-plane operation.

The Reader and Data Access role allows a user to:

* View storage account properties (Reader)
* Read data within storage services such as blobs or files (Data Access) However, according to Microsoft Azure RBAC documentation, this role does not include permissions to list or regenerate storage account keys. Key management operations require permissions such as:
* Microsoft.Storage/storageAccounts/listKeys/action
* Microsoft.Storage/storageAccounts/regenerateKey/action

These permissions are included in roles such as:

* Storage Account Contributor
* Contributor
* Owner

Because the Reader and Data Access role lacks key management permissions, assigning this role to User1 does not allow them to list or regenerate storage account keys for storage1.

Microsoft explicitly documents that only management roles, not data access roles, can manage storage account keys. Therefore, the proposed solution does not meet the goal.

NEW QUESTION: 167

You have a virtual network named VNet1 as shown in the exhibit. (Click the Exhibit tab.)

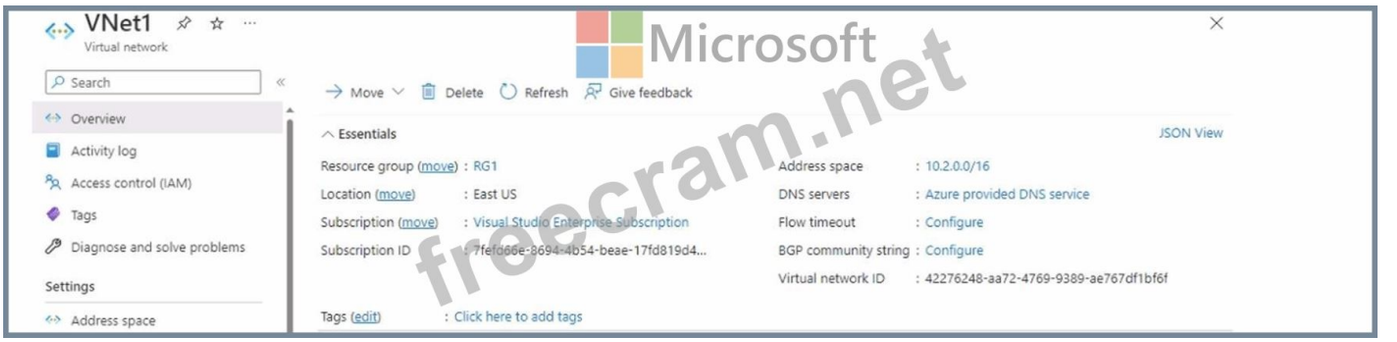

No devices are connected to VNet1.

You plan to peer VNet1 to another virtual network named VNet2. VNet2 has an address space of 10.2.0.0/16.

You need to create the peering.

What should you do first?

A. Configure a service endpoint on VNet2.

B. Modify the address space of VNet1.

C. Add a gateway subnet to VNet1.

D. Create a subnet on VNet1 and VNet2.

Answer: (SHOW ANSWER)

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: 168

You have an Azure subscription that contains the storage accounts shown in the following table.

You plan to use AzCopy to copy a blob from contained directly to share. You need to identify which authentication method to use when you use AzCopy.

What should you identify for each account? To answer, drag the appropriate authentication methods to the correct accounts. Each method may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:
storage1: A shared access signature (SAS) token
storage2: A shared access signature (SAS) token

NEW QUESTION: 169
You have an Azure subscription that contains two peered virtual networks named VNet1 and VNet2. VNet1 has a VPN gateway that uses static routing.
The on-premises network has a VPN connection that uses the VPN gateway of VNet1.
You need to configure access for users on the on-premises network to connect to a virtual machine on VNet2.
The solution must minimize costs.
Which type of connectivity should you use?
A. Azure Firewall with a private IP address
B. ExpressRoute circuits to VNet2
C. service chaining and user-defined routes (UDRs)
D. Azure Application Gateway
Answer: (SHOW ANSWER)
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: 170
You need to implement the planned changes for the new containers.
Which Azure services can you use for each image? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Answer:

Answer Area

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: 171
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an Azure virtual machine named VM1. VM1 was deployed by using a custom Azure Resource Manager template named ARM1.json.
You receive a notification that VM1 will be affected by maintenance.
You need to move VM1 to a different host immediately.
Solution: From the Update management blade, you click Enable.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
Azure virtual machines can occasionally be subject to platform maintenance events. Maintenance

can be planned (with notice) or unplanned. Planned maintenance includes updates to the host operating system or hardware that may require moving your VM to another host.
When you receive a maintenance notification, Azure provides options to move your VM proactively to minimize downtime. According to the Microsoft Azure Administrator documentation, the correct action is to use the "Redeploy" option or "Move VM to another host" from the Azure portal. This operation shuts down the VM, migrates it to a new physical node within the same region, and then restarts it - effectively resolving maintenance impact.
The Update management blade, however, is part of Azure Automation Update Management, which is used to schedule and manage operating system updates (Windows or Linux) across Azure and hybrid environments.
Enabling Update Management simply connects the VM to a Log Analytics workspace and does not move or redeploy the VM.
Therefore, enabling Update Management does not achieve the goal of moving VM1 to a different host to avoid maintenance. The correct action would be:
Navigate to VM1 # Help # Redeploy + reapply # Redeploy.
This redeploys the VM to another host server while retaining its network and disk configurations.
Thus, the proposed solution does not meet the goal.

NEW QUESTION: 172
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You manage a virtual network named VNet1 that is hosted in the West US Azure region.
VNet1 hosts two virtual machines named VM1 and VM2 that run Windows Server.
You need to inspect all the network traffic from VM1 to VM2 for a period of three hours.
Solution: From Performance Monitor, you create a Data Collector Set (DCS).
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
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: 173
You sign up for Azure Active Directory (Azure AD) Premium.
You need to add a user named admin1@contoso.com as an administrator on all the computers that will be joined to the Azure AD domain.
What should you configure in Azure AD?
A. Device settings from the Devices blade.
B. General settings from the Groups blade.
C. User settings from the Users blade.
D. Providers from the MFA Server blade.
Answer: (SHOW ANSWER)
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: 174
Which blade should you instruct the finance department auditors to use?
A. Partner information
B. Overview
C. Payment methods
D. Invoices
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: 175
You have an Azure Kubernetes Service (AKS) cluster named AKS1.
You need to configure cluster autoscaler for AKS1.
Which two tools should you use? Each correct answer presents a complete solution, NOTE: Each correct selection is worth one point
A. the set-AzAKs cmdlet
B. the Azure portal
C. The az aks command
D. the kubect1 command
E. the set Azure cmdlet
Answer: (SHOW ANSWER)
AKS clusters can scale in one of two ways: - The cluster autoscaler watches for pods that can't be scheduled on nodes because of resource constraints. The cluster then automatically increases the number of nodes. - The horizontal pod autoscaler uses the Metrics Server in a Kubernetes cluster to monitor the resource demand of pods. If an application needs more resources, the number of pods is automatically increased to meet the demand. Reference: https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler
NEW QUESTION: 176
You have an Azure subscription that contains the resources shown in the following table.

Name	Resource group	Type	Location
app1	RG1	Container app	East US
Vault1	RG1	Azure Key Vault	East US
Vault2	RG1	Azure Key Vault	West US
Vault3	RG2	Azure Key Vault	East US

You plan to use an Azure key vault to provide a secret to appl.
What should you create for app1 to access the key vault, and from which key vault can the secret be used? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

NEW QUESTION: 177
You have an Azure subscription that contains a storage account named storageacct1234 and two users named User1 and User2.
You assign User1 the roles shown in the following exhibit.

Which two actions can User1 perform? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. View file shares in storageacct1234.

B. Upload blob data to storageacct1234.

C. Assign roles to User2 for storageacctl234.

D. View blob data in storageacctl234.

E. Modify the firewall of storageacct1234.

Answer: (SHOW ANSWER)

n this scenario, User1 has two role assignments:

* Reader (Inherited from Resource Group scope)
* Scope: Resource Group (inherited)
* Permission type: Read-only access
* Allows User1 to view the configuration and properties of Azure resources within the resource group (including the storage account).
* However, it does not allow data operations, such as reading or writing blob contents.
* Storage Blob Data Contributor (Assigned at "This resource" level)
* Scope: storageacct1234 (this resource)
* Permission type: Data access for blob content
* Allows read, write, and delete operations on blob data in the storage account.

According to Microsoft Azure built-in roles documentation:

* Reader role: "View Azure resources but cannot make any changes."
* Storage Blob Data Contributor role: "Grants read, write, and delete permissions to Azure Storage blob containers and data." This means:
* User1 can upload blobs because the Storage Blob Data Contributor role provides write permissions.
* User1 can view blob data because the same role provides read access to blob content.
* User1 cannot assign roles (no RBAC management rights, as the role doesn't include Microsoft. Authorization/* actions).

* User1 cannot modify the firewall or network settings of the storage account - those actions require Contributor or Owner roles at the management plane level.
* User1 can view file shares under the storage account metadata due to Reader access but cannot access data within them because that requires Storage File Data SMB Share Reader/Contributor permissions (different data plane).
# Final Verified Answers:
* B. Upload blob data to storageacct1234
* D. View blob data in storageacct1234
Summary (from Microsoft Role Documentation Extracts):
* Storage Blob Data Contributor = Can read/write/delete blob data (Data Plane)
* Reader = Can view Azure resource configuration (Management Plane)
* No permission to assign roles or modify security/firewall settings.

NEW QUESTION: 178
You need to ensure that User1 can create initiative definitions, and User4 can assign initiatives to RG2. The solution must meet the technical requirements.
Which role should you assign to each user? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

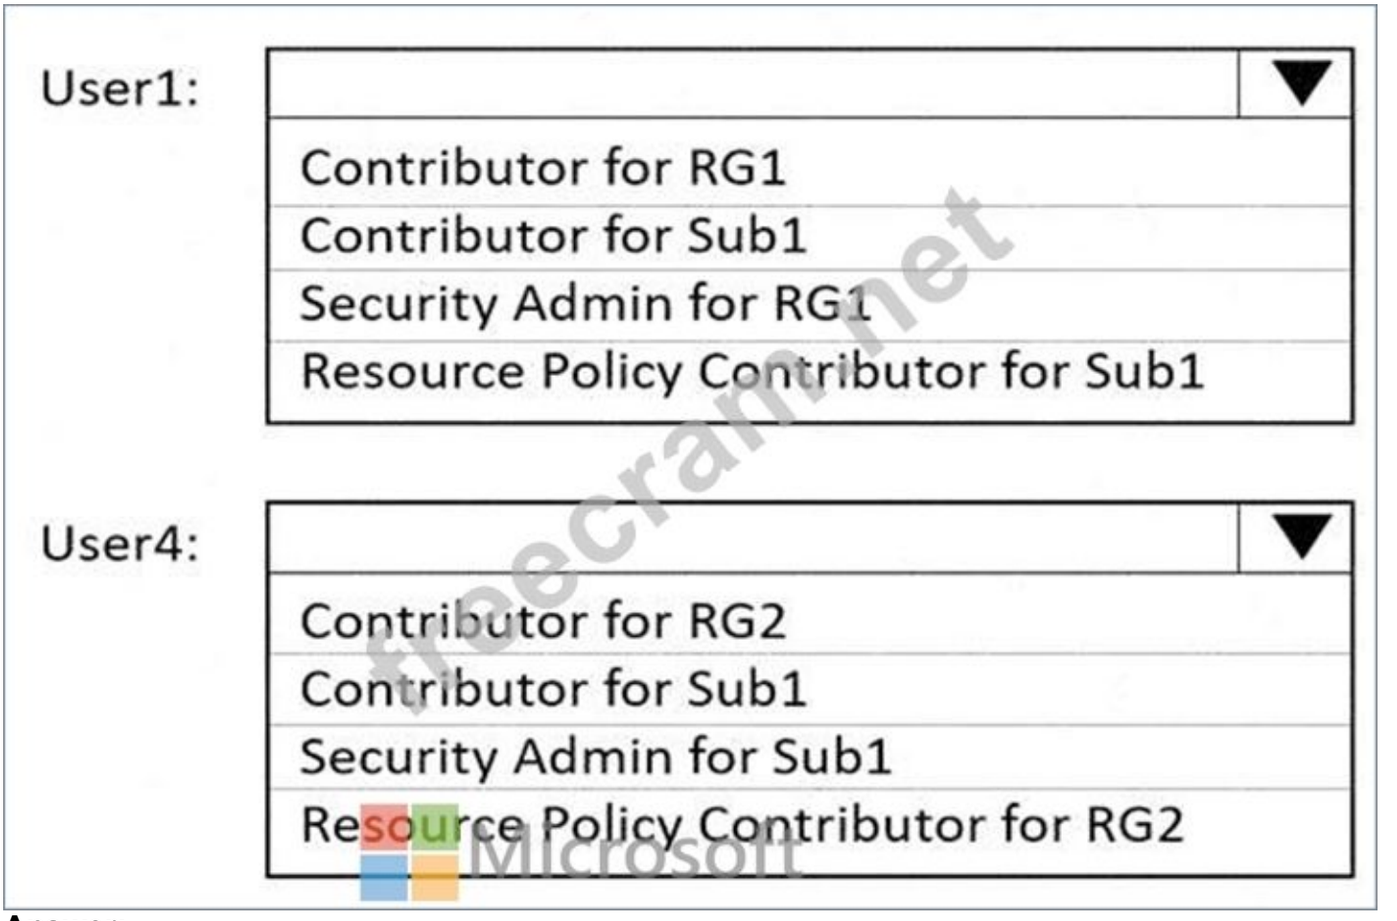

Answer:

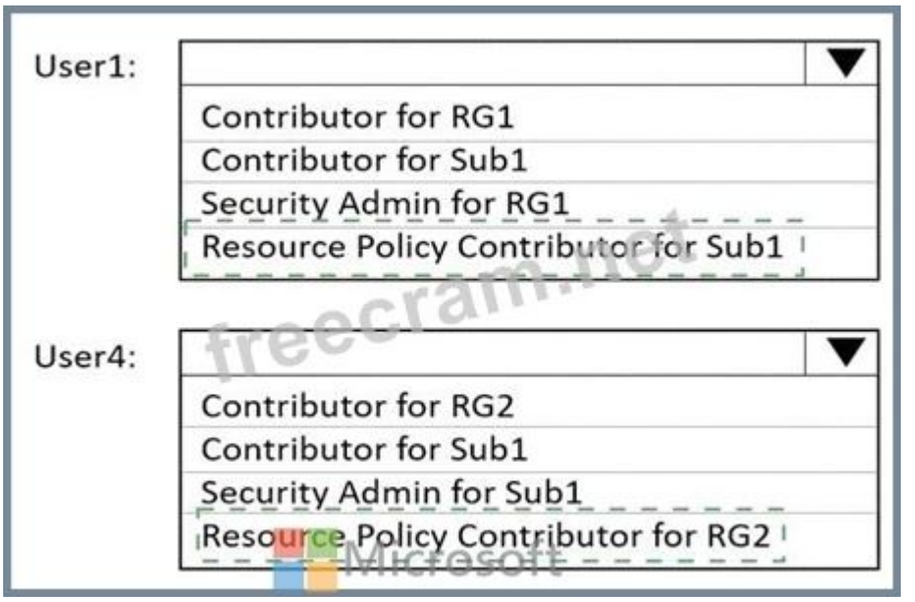

Explanation:

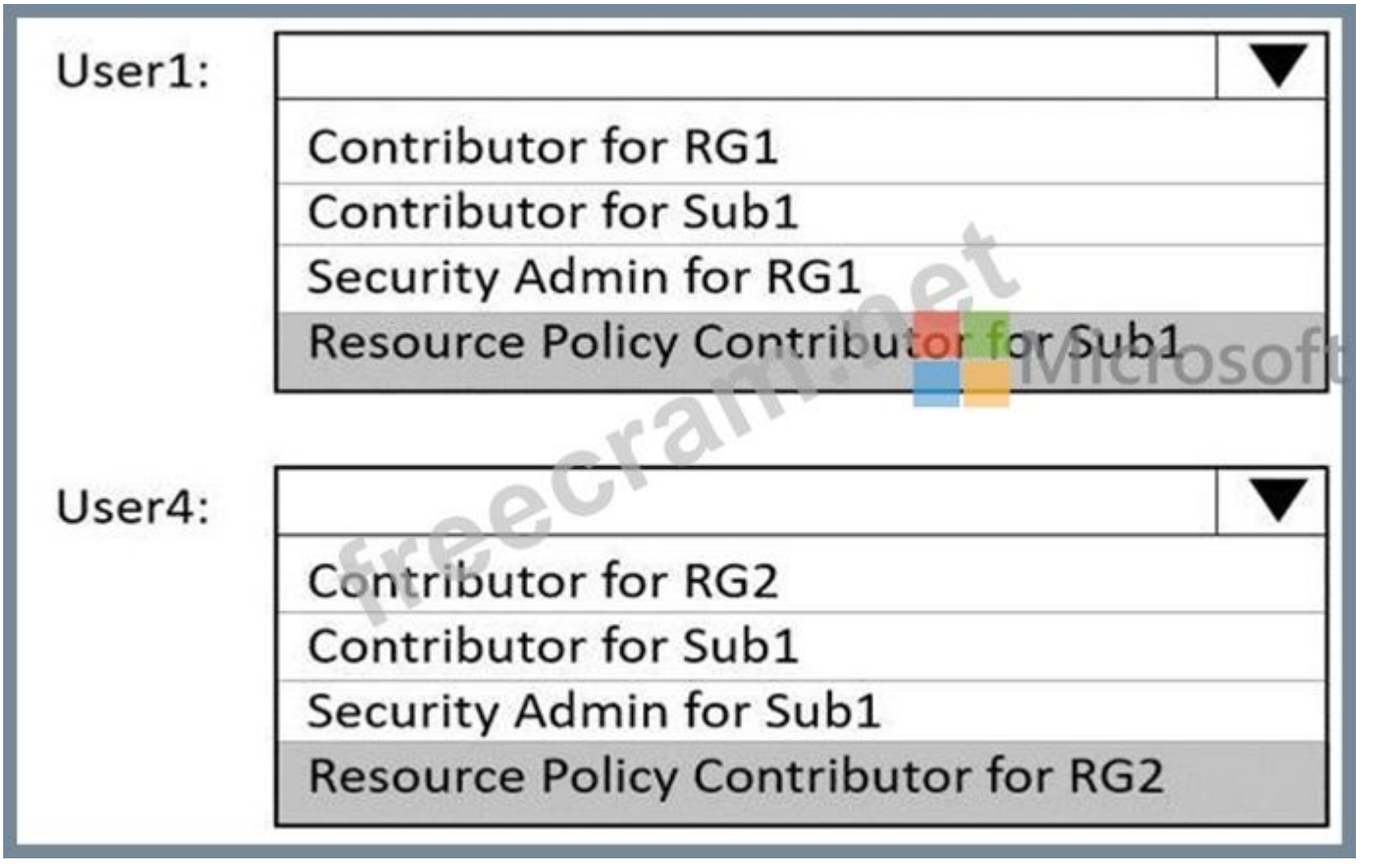

To meet the technical requirement that User1 can create initiative definitions and User4 can assign initiatives to RG2, we must understand how Azure Policy roles operate according to Microsoft Azure documentation and AZ-104 study guides.

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: 179
You have a hybrid deployment of Azure AD that contains the users shown in the following table.

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

You need to modify the JobTitle and UsageLocation attributes for the users.
For which users can you modify the attributes from Azure AD? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

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: 180
You are evaluating the connectivity between the virtual machines after the planned implementation of the Azure networking infrastructure.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Statements	Yes	No
The virtual machines on Subnet1 will be able to connect to the virtual machines on Subnet3.	○	○
The virtual machines on ClientSubnet will be able to connect to the Internet.	○	○
The virtual machines on Subnet3 and Subnet4 will be able to connect to the Internet.	○	○

Answer:

Statements	Yes	No
The virtual machines on Subnet1 will be able to connect to the virtual machines on Subnet3.	○	○
The virtual machines on ClientSubnet will be able to connect to the Internet.	○	○
The virtual machines on Subnet3 and Subnet4 will be able to connect to the Internet.	○	○

Explanation:

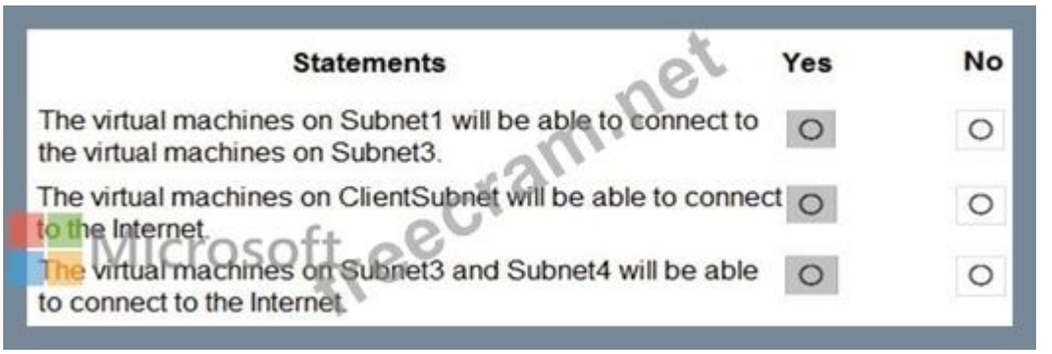

Statements	Yes	No
The virtual machines on Subnet1 will be able to connect to the virtual machines on Subnet3.	○	○
The virtual machines on ClientSubnet will be able to connect to the Internet.	○	○
The virtual machines on Subnet3 and Subnet4 will be able to connect to the Internet.	○	○

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: 181
You have an Azure Active Directory (Azure AD) tenant named adatum.com. Adatum.com contains the groups in the following table.

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*

You create two user accounts that are configured as shown in the following table.

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

To which groups do User1 and User2 belong? To answer. select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

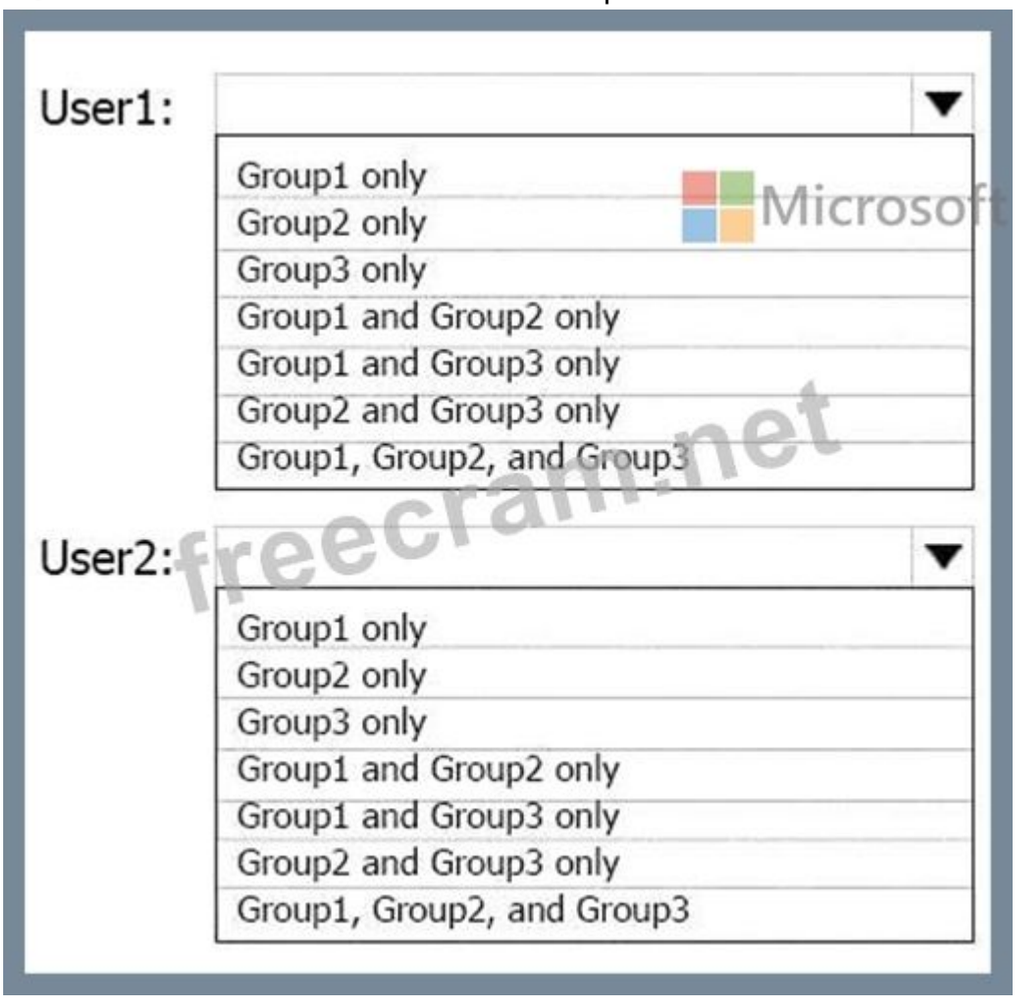

Answer:

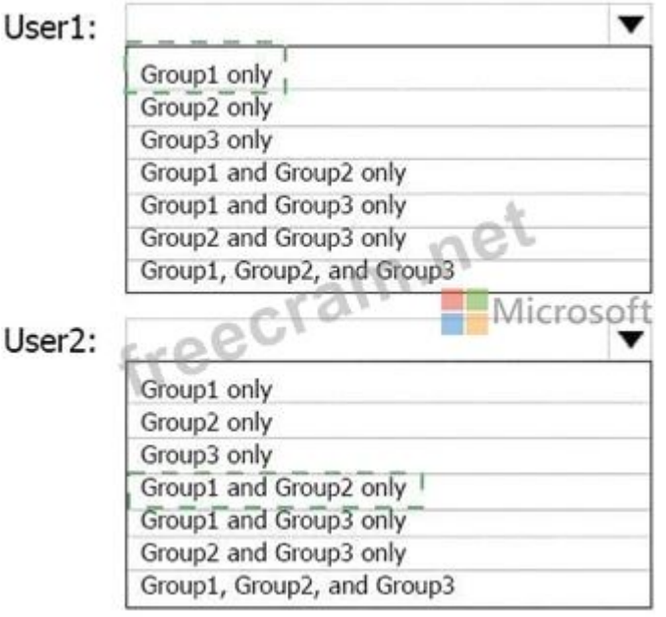

Explanation:

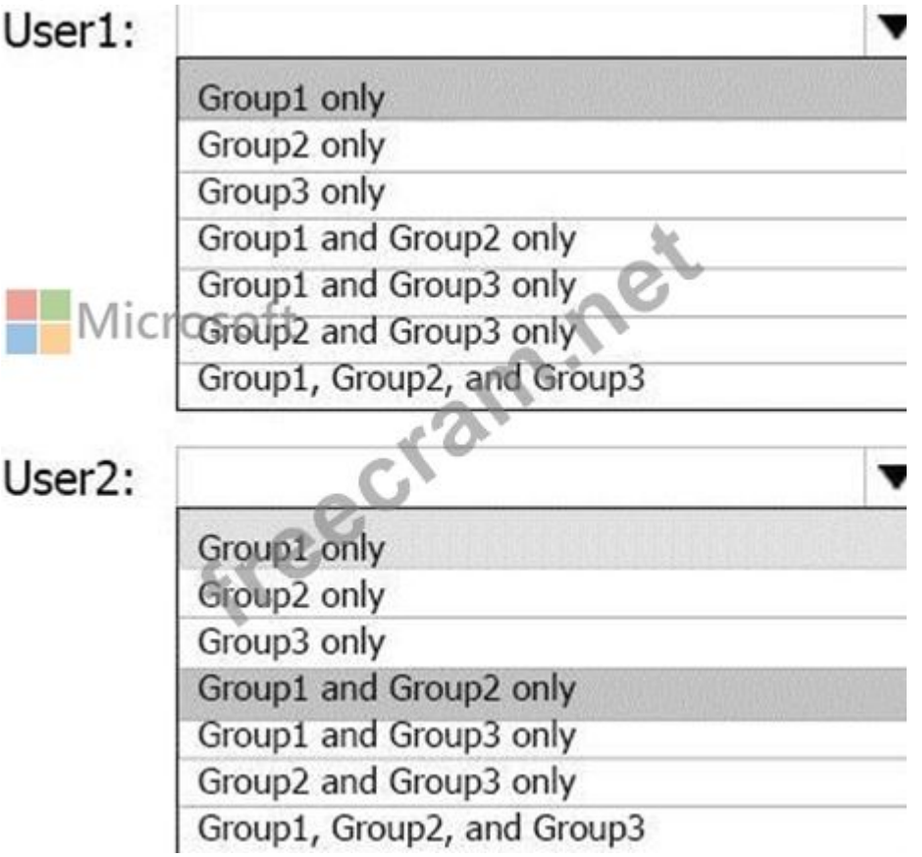

Azure Active Directory (Azure AD) supports three types of groups:
Security groups (for access and permissions),
Microsoft 365 groups (for collaboration tools such as Teams, SharePoint, and Outlook), Dynamic groups (with automatically managed memberships based on user attributes).
A dynamic user group 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: 182
You have an on-premises network.
You have an Azure subscription that contains two virtual networks named VNet1 and VNet2. VNet1 uses an IP address space of 192.168.8.0/24. VNet2 uses an IP address space of 192.168.9.0/24.
You need to configure the virtual networks. The solution must meet the following requirements:
* Ensure that the resources on VNet1 can communicate with the resources on VNet2.
* Ensure that the resources on the on-premises network can communicate with Azure resources.
* Minimize costs.

Answer:

Explanation:

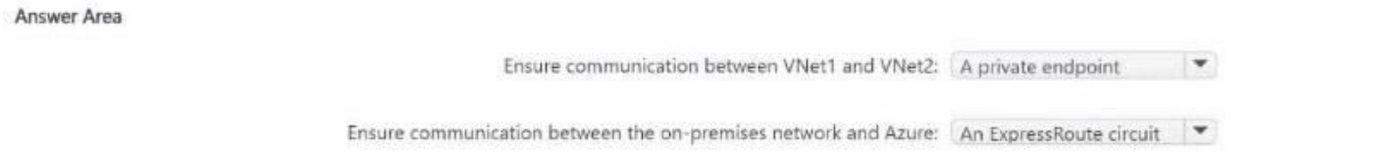

To allow VNet1 (192.168.8.0/24) and VNet2 (192.168.9.0/24) to communicate, the lowest-overhead Azure- native option is virtual network peering. Azure documentation describes VNet peering as a way to connect two VNets so that resources can communicate using private IP addresses as if they were on the same network.

Because the address spaces do not overlap, peering is supported and does not require deploying extra appliances. This also minimizes administrative effort compared to routing through gateways or NVAs.

For on-premises to Azure connectivity while minimizing cost, the appropriate service is an Azure VPN Gateway with a site-to-site (S2S) VPN. Azure's guidance for hybrid connectivity positions VPN Gateway as the cost-effective encrypted tunnel over the public internet for connecting on-premises networks to Azure VNets. The alternative shown, ExpressRoute, provides private dedicated connectivity but is designed for higher throughput/enterprise requirements and generally increases cost due to circuit provisioning and associated charges. Therefore, combining VNet peering (VNet-to-VNet) with VPN Gateway (on-prem-to- Azure) satisfies both communication requirements and the cost constraint.

NEW QUESTION: 183

You have an Azure subscription that contains the users shown in the following table.

The groups are configured as shown in the following table.

Name	Type	Azure AD roles can be assigned to the group
Group1	Security	Yes
Group2	Security	Yes
Group3	Microsoft 365	Yes

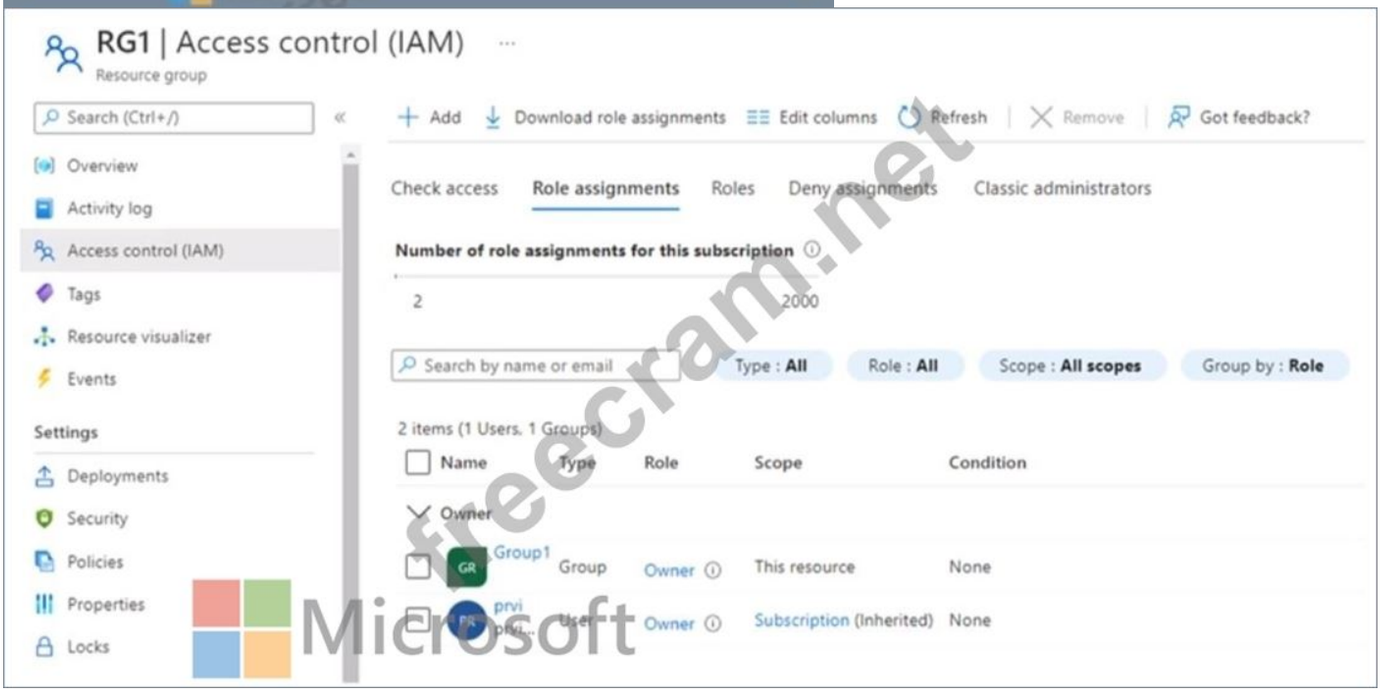

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
You can assign User2 the Owner role for RG1 by adding Group2 as a member of Group1.	○	○
You can assign User3 the Owner role for RG1 by adding Group3 as a member of Group1.	○	○
You can assign User3 the Owner role for RG1 by assigning the Owner role to Group3 for RG1.	○	○

Answer:

Answer Area

Statements	Yes	No
You can assign User2 the Owner role for RG1 by adding Group2 as a member of Group1.	○	○
You can assign User3 the Owner role for RG1 by adding Group3 as a member of Group1.	○	○
You can assign User3 the Owner role for RG1 by assigning the Owner role to Group3 for RG1.	◉	○

Explanation:

Answer Area

Statements	Yes	No
You can assign User2 the Owner role for RG1 by adding Group2 as a member of Group1.	○	◉
You can assign User3 the Owner role for RG1 by adding Group3 as a member of Group1.	○	◉
You can assign User3 the Owner role for RG1 by assigning the Owner role to Group3 for RG1.	◉	○

https://learn.microsoft.com/en-us/azure/active-directory/roles/groups-concept#how-are-role-

assignable- groups-protected
"Group nesting isn't supported. A group can't be added as a member of a role-assignable group."
For the second question:
https://learn.microsoft.com/en-us/azure/active-directory/fundamentals/how-to-manage-groups#add-or-remove- a-group-from-another-group
"We currently don't support:
Adding Microsoft 365 groups to Security groups or other Microsoft 365 groups.
"
For the third question, although it appears truncated in the screenshot (ending with "for...") there is a reference about Microsoft 365 groups support for roles assignment here:
https://learn.microsoft.com/en-us/azure/active-directory/roles/groups-concept#how-role-assignments-to- groups-work
"To assign a role to a group, you must create a new security or Microsoft 365 group with the isAssignableToRole property set to true. "

NEW QUESTION: 184
You have an Azure subscription that contains a storage account named storage 1.
You need to ensure that the access keys for storage! rotate automatically.
What should you configure?
A. a backup vault
B. redundancy for storage1
C. lifecycle management for storage1
D. an Azure key vault
E. a Recovery Services vault
Answer: (SHOW ANSWER)
In Azure, a storage account access key provides full access to all data within the account. To reduce risk and follow security best practices, these keys should be rotated (regenerated) periodically.
According to the Microsoft Azure Storage and Security documentation, Azure Key Vault can be used to automate access key rotation for storage accounts.
Azure Key Vault allows you to:
* Store and manage secrets, keys, and certificates securely.
* Integrate directly with Azure Storage to manage account keys and Shared Access Signatures (SAS).
* Enable automated key rotation when using Azure Key Vault managed storage account keys.
Here's how it works:
* In Azure Key Vault, add the storage account (storage1) as a managed storage account.
* Key Vault periodically regenerates (rotates) the storage access keys automatically.
* Applications can retrieve updated keys via Key Vault APIs or managed identities without manual key updates.
This process ensures consistent security and reduces the administrative effort required for key

rotation.
Other options such as backup vaults, redundancy, or lifecycle management do not handle access key rotation-they serve data protection or retention purposes, not key management.
Final Verified Answer: D. an Azure key vault

NEW QUESTION: 185
You have an Azure subscription that contains two Log Analytics workspaces named Workspace 1 and Workspace? and 100 virtual machines that run Windows Server.
You need to collect performance data and events from the virtual machines. The solution must meet the following requirements:
* Logs must be sent to Workspace! and Workspace?
* All Windows events must be captured
* All security events must be captured.
What should you install and configure on each virtual machine?
A. the Azure Monitor agent
B. the Windows Azure diagnostics extension (WAD)
C. the Windows VM agent
Answer: (SHOW ANSWER)
https://learn.microsoft.com/en-us/azure/azure-monitor/agents/agents-overview Azure Monitor Agent (AMA) collects monitoring data from the guest operating system of Azure and hybrid virtual machines and delivers it to Azure Monitor for use by features, insights, and other services, such as Microsoft Sentinel and Microsoft Defender for Cloud. Azure Monitor Agent replaces all of Azure Monitor's legacy monitoring agents.

NEW QUESTION: 186
You have an Azure subscription named Subscription1. Subscription1 contains two Azure virtual machines named VM1 and VM2. VM1 and VM2 run Windows Server 2016.
VM1 is backed up daily by Azure Backup without using the Azure Backup agent.
VM1 is affected by ransomware that encrypts data.
You need to restore the latest backup of VM1.
To which location can you restore the backup? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

You can perform a file recovery of VM1 to:

- VM1 only
- VM1 or a new Azure virtual machine only
- VM1 and VM2 only
- A new Azure virtual machine only
- Any Windows computer that has Internet connectivity

You can restore VM1 to:

- VM1 only
- VM1 or a new Azure virtual machine only
- VM1 and VM2 only
- Any Windows computer that has Internet connectivity

Answer:

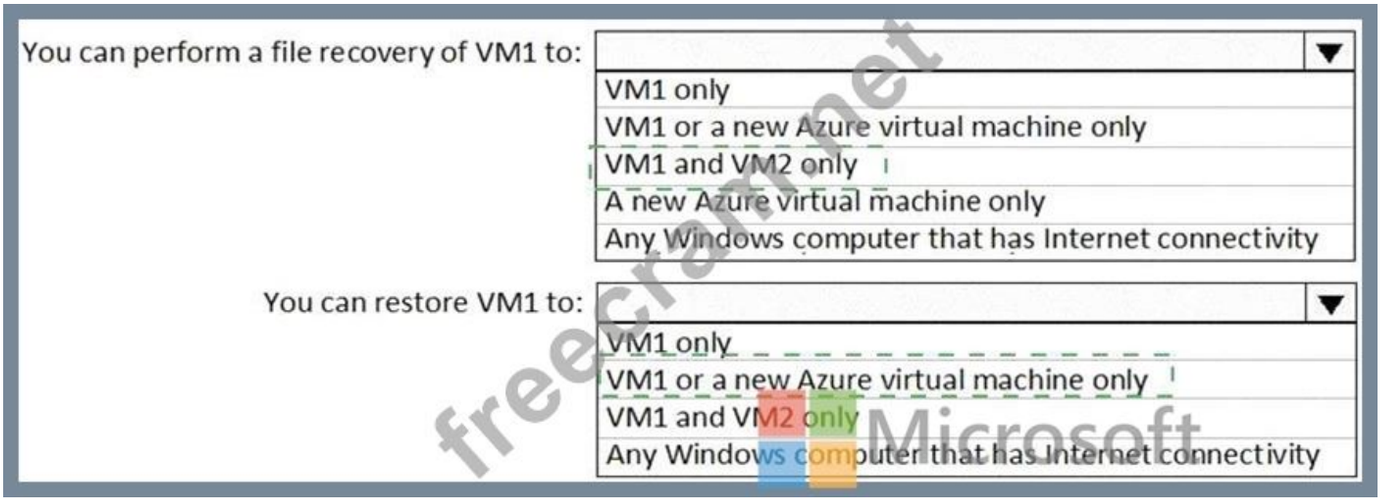

Explanation:

Box 1 : VM1 and VM2 only

When recovering files, you can't restore files to a previous or future operating system version.You can restore files from a VM to the same server operating system, or to the compatible client operating system. Therefore -

"VM1 and VM2 only" is the best answer since both run on Windows Server 2016.

"A new Azure virtual machine only" ,this will also work but why to create unnecessary new VM in Azure if existing VM will do the task. So this option is incorrect.

Box 2 : VM1 or A new Azure virtual machine only

When restoring a VM, you can't use the replace existing VM option for encrypted VMs. This option is only supported for unencrypted managed disks. And also You can restore files from a VM to the same server operating system, or to the compatible client operating system only.

Hence "VM1 or A new Azure virtual machine only" is correct answer.

Answer Area

You can perform a file recovery of VM1 to:

- VM1 only
- VM1 or a new Azure virtual machine only
- **VM1 and VM2 only**
- A new Azure virtual machine only
- Any Windows computer that has Internet connectivity

You can restore VM1 to:

- VM1 only
- **VM1 or a new Azure virtual machine only**
- VM1 and VM2 only
- Any Windows computer that has Internet connectivity

References:
https://docs.microsoft.com/en-us/azure/backup/backup-azure-arm-restore-vms
https://docs.microsoft.com/en-us/azure/backup/backup-azure-restore-files-from-vm#system-requirements

NEW QUESTION: 187
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Resource group	Region
Vault1	Recovery services vault	RG1	East US
VM1	Virtual machine	RG1	East US
VM2	Virtual machine	RG1	West US

All virtual machines run Windows Server.
On VM1, you back up a folder named Folder! as shown in the following exhibit

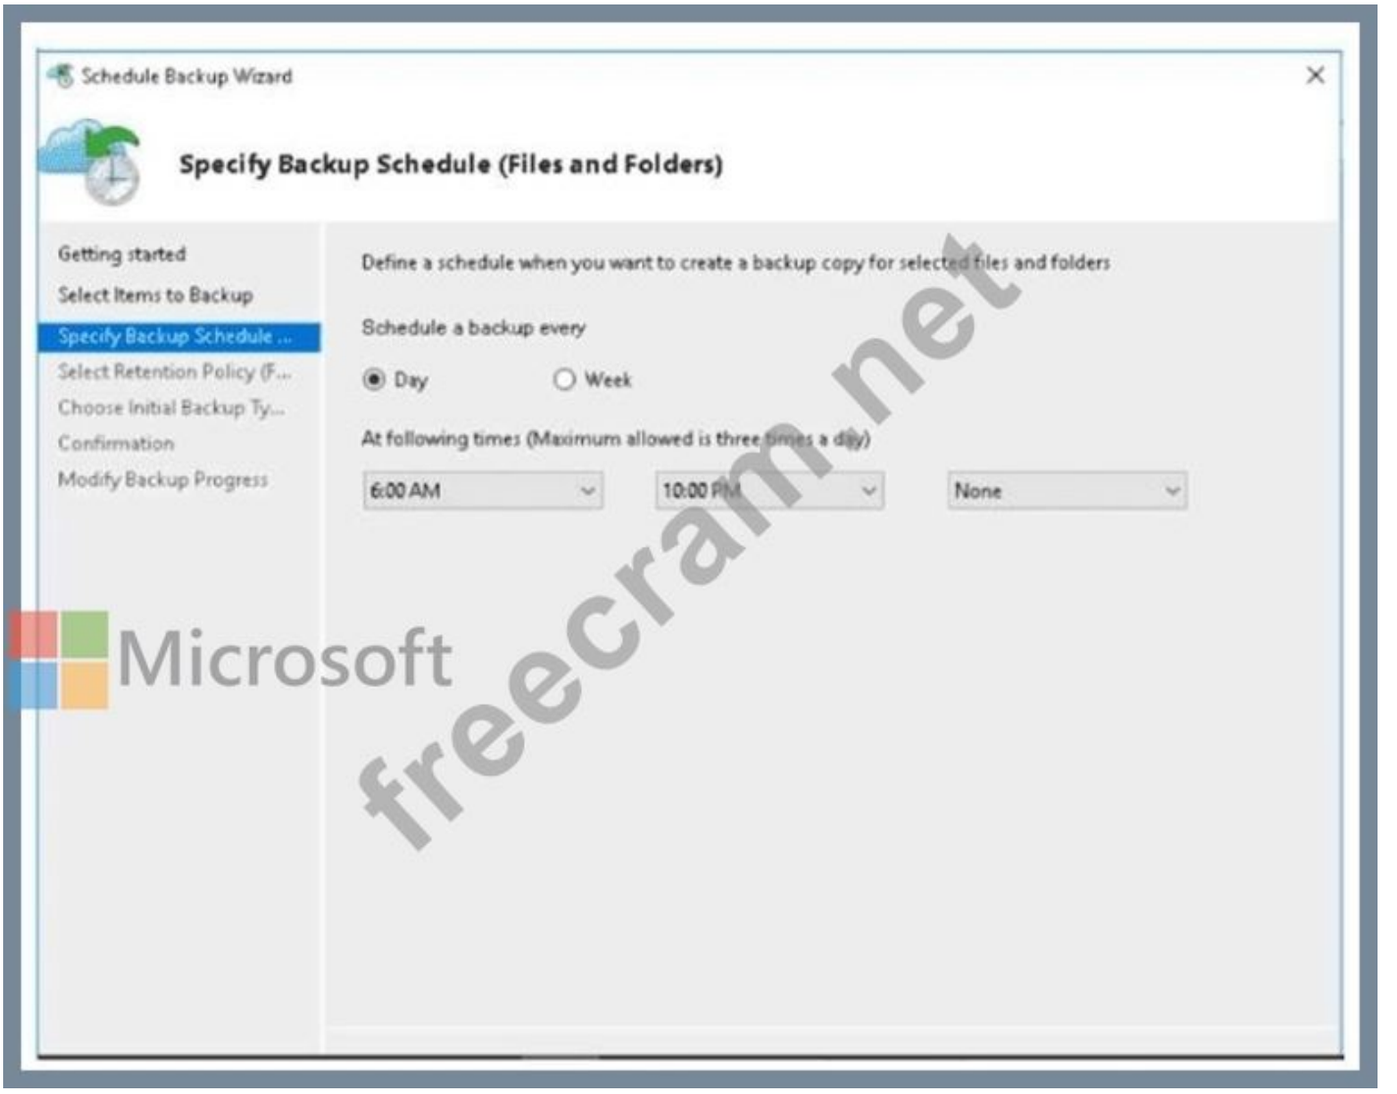

You plan to restore the backup to a different virtual machine.

You need to restore the backup to VM2.

What should you do first?

A. From VM1, install the Microsoft Azure Recovery Services Agent

B. From VM1, install the Windows Server Backup feature.

C. From VM2, install the Windows Server Backup feature.

D. From VM2, install the Microsoft Azure Recovery Services Agent.

Answer: (SHOW ANSWER)

The screenshot shows the Schedule Backup Wizard (Files and Folders), which is the Microsoft Azure Recovery Services (MARS) agent experience. This is file/folder backup to a Recovery Services vault, not an

"Azure VM backup" policy-based backup. In Microsoft's Azure Backup guidance, the MARS agent protects files and folders on a Windows machine by uploading encrypted backup data to the Recovery Services vault.

The encryption is controlled by a customer-managed passphrase, and Azure does not store that passphrase.

Because MARS backups are agent-based, restore is not limited to the original VM or to the VM's region.

Microsoft documentation for MARS restore describes restoring to the original location or to an alternate location, including another server, as long as the target machine is registered to the same vault and you provide the same passphrase used to encrypt the backups. The presence of VM2 in West US does not block restore for MARS-protected files/folders; region coupling is a constraint for Azure VM backup (vault and VM must be in the same region), not for MARS agent file/folder backup.
So, after redeploying or on another VM, you install and register the MARS agent to Vault1, then perform a restore of Folder1 using the correct passphrase-VM2 is a valid restore target.

NEW QUESTION: 188
You manage two Azure subscriptions named Subscription 1 and Subscription2.
Subscription! has following virtual networks:

Name	Address space	Region
VNET1	10.10.10.0/24	West Europe
VNET2	172.16.0.0/16	West US

The virtual networks contain the following subnets:

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 contains the following virtual network:
- Name: VNETA
* Address space: 10.10.128.0/17
* Region: Canada Central
VNETA contains the following subnets:

Name	Address range
SubnetA1	10.10.130.0/24
SubnetA2	10.10.131.0/24

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

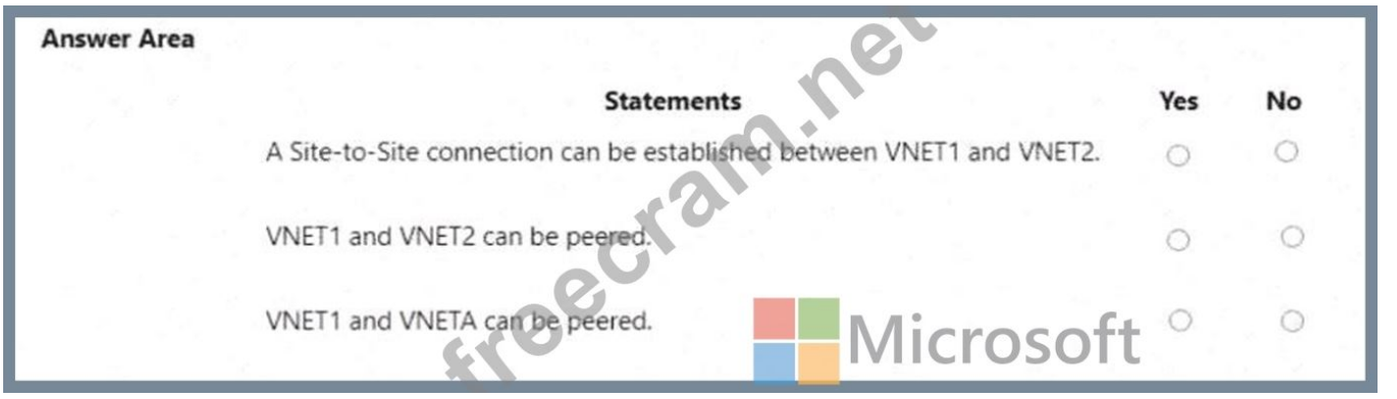

Answer:

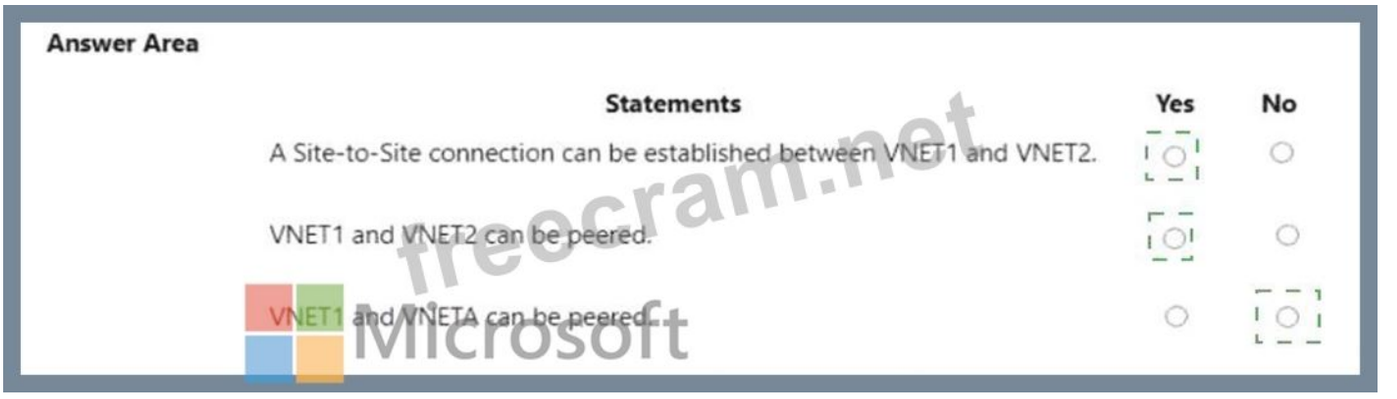

Explanation:

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.

NEW QUESTION: 189
You have an Azure subscription.
You plan to deploy a container.
You need to recommend which Azure services can scale the container automatically.
What should you recommend?
A. Azure Container Apps only
B. Azure Container Instances only
C. Azure Container Apps or Azure App Service only
D. Azure Container Instances or Azure App Service only
E. Azure Container Apps, Azure Container Instances, or Azure App Service
Answer: (SHOW ANSWER)
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: 190
You have an Azure AD tenant that is linked to the subscriptions shown in the following table.

Name	Management group	Parent management group
Sub1	Tenant Root Group	*Not applicable*
Sub2	MG1	Tenant Root Group
Sub3	MG2	Tenant Root Group

You have the resource groups shown In the following table.

Name	Subscription	Description
RG1	Sub1	Contains a storage account named storage1
RG2	Sub2	Contains a web app named App1
RG3	Sub3	Contains a virtual machine named VM1

You assign roles to users as shown in the following table.

User	Role	Scope
User1	Contributor	MG2
User2	Storage Account Contributor	storage1
User3	User Access Administrator	Tenant Root Group

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:
User1 can resize VM1. Yes, this is correct. According to the tables, User1 is assigned the Contributor role at the subscription level for Sub1. The Contributor role grants full access to manage all resources in the subscription, including the ability to resize virtual machines1. Therefore, User1 can resize VM1, which is a resource in RG1 under Sub1.
User2 can create a new storage account in RG1. No, this is not correct. According to the tables, User2 is assigned the Reader role at the resource group level for RG1. The Reader role grants read-only access to view existing resources in the resource group, but not to create, update, or delete any resources2. Therefore, User2 cannot create a new storage account in RG1.
User3 can assign User1 the Owner role for RG3. No, this is not correct. According to the tables, User3 is assigned the Storage Account Contributor role at the resource group level for RG3. The Storage Account Contributor role grants full access to manage storage accounts and their data in the resource group, but not to assign roles to other users3. To assign roles to other users, User3 would need a role that has Microsoft.
Authorization/roleAssignments/write permissions, such as User Access Administrator or Owner4.

Therefore, User3 cannot assign User1 the Owner role for RG3.

NEW QUESTION: 191
You create an Azure Storage account.
You plan to add 10 blob containers to the storage account.
For one of the containers, you need to use a different key to encrypt data at rest.
What should you do before you create the container?
A. Modify the minimum TLS version.
B. Create an encryption scope.
C. Generate a shared access signature (SAS).
D. Rotate the access keys.
Answer: (SHOW ANSWER)
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: 192
You have an Azure subscription that contains a virtual machine named VM1.
You have an on-premises datacenter that contains a domain controller named DC1.
ExpressRoute is used to connect the on-premises datacenter to Azure.
You need to use Connection Monitor to identify network latency between VM1 and DC1.
What should you install on DC1?
A. the Log Analytics agent
B. the Azure Network Watcher Agent virtual machine extension
C. an Azure Monitor agent extension
D. the Azure Connected Machine agent for Azure Arc-enabled servers

Answer: (SHOW ANSWER)

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: 193
You have an app named App1 that runs on an Azure web app named webapp1.
The developers at your company upload an update of App1 to a Git repository named GUI.
Webapp1 has the deployment slots shown in the following table.

Name	Function
webapp1-prod	Production
webapp1-test	Staging

You need to ensure that the App1 update is tested before the update is made available to users.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.
A. Swap the slots
B. Deploy the App1 update to webapp1-prod, and then test the update
C. Stop webapp1-prod
D. Deploy the App1 update to webapp1-test, and then test the update
E. Stop webapp1-test
Answer: (SHOW ANSWER)
Azure App Service Deployment Slots allow developers to deploy new versions of applications to a staging environment (for example, webapp1-test) before pushing them to production (webapp1-prod).
The Azure Administrator documentation states:
"You can deploy your app to a non-production slot, test it, and then swap it into production with zero downtime. Swapping exchanges the content and configuration of the two slots, making the

tested app live." Here's the correct process:
1## Deploy the App1 update to the staging slot (webapp1-test) - This allows testing of the new version without impacting production users.
2## Test the application in the staging slot to ensure all functionality is correct.
3## Swap the slots (Swap webapp1-test with webapp1-prod) - This promotes the tested code to production with no downtime, and the old production version moves to the staging slot automatically.
This workflow ensures quality validation before deployment and minimal production interruption.
Final Verified Answers for Question 284:
* A. Swap the slots
* D. Deploy the App1 update to webapp1-test, and then test the update

NEW QUESTION: 194
You have an Azure Storage account named storage1.
You need to enable a user named User1 to list and regenerate storage account keys for storage1.
Solution: You assign the Storage Account Key Operator Service Role to User1.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In Azure Role-Based Access Control (RBAC), the Storage Account Key Operator Service Role allows a user to list and regenerate access keys for a storage account but does not grant permission to access the data itself.
The Storage Account Key Operator Service Role has the following key permissions:
* Microsoft.Storage/storageAccounts/listkeys/action
* Microsoft.Storage/storageAccounts/regeneratekey/action
* Microsoft.Storage/storageAccounts/read
This role is typically used by automation or service accounts that need to manage storage account keys without having full access to the data stored within the account (blobs, files, queues, etc.).
In the scenario, User1 only needs to list and regenerate storage account keys for storage1, not perform data operations or modify configurations. The Storage Account Key Operator Service Role provides precisely these permissions, making it the correct and minimal-privilege choice for this task.
Thus, assigning this role to User1 meets the goal as per Azure RBAC guidelines and Microsoft's least privilege principle documented in the Azure Administrator exam guide.
NEW QUESTION: 195
You have an Azure subscription that contains eight virtual machines and the resources shown in the following table.

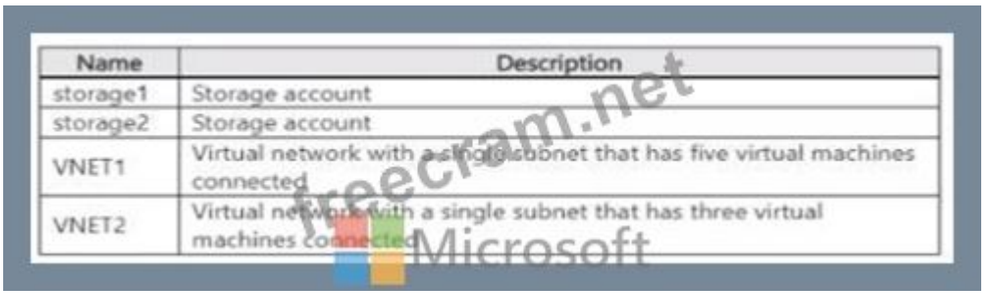

Name	Description
storage1	Storage account
storage2	Storage account
VNET1	Virtual network with a single subnet that has five virtual machines connected
VNET2	Virtual network with a single subnet that has three virtual machines connected

You need to configure access for VNEI1. The solution must meet the following requirements:
* The virtual machines connected to VNET1 must be able to communicate with the virtual machines connected to VNET2 by using the Microsoft backbone.
* The virtual machines connected to VNETI must be able to access storage1. Storage2, and Microsoft Entra ID by using the Microsoft backbone.
What is the minimum number of service endpoints you should add to VNET1?
A. 1
B. 2
C. 3
D. 5
Answer: (SHOW ANSWER)
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: 196
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have a Microsoft Entra tenant named Adatum.com and an Azure Subscription named Subscription1.
Adatum.com contains a group named Developers. Subscription1 contains a resource group named Dev.

You need to provide the Developers group with the ability to create Azure logic apps in the Dev resource group.
Solution: On Dev, you assign the Logic App Contributor role to the Developers group.
Does this meet the goal?
A. Yes
B. No
Answer: A (LEAVE A REPLY)
In Microsoft Azure, Role-Based Access Control (RBAC) allows you to manage access to Azure resources by assigning roles to users, groups, and service principals. The Logic App Contributor role specifically grants the ability to create, edit, and manage Logic Apps but not assign permissions or modify access control (IAM).
According to the official Microsoft Azure documentation for built-in roles in Azure RBAC:
"The Logic App Contributor role lets you manage logic apps, but not access them. You can view, edit, update, and delete logic apps." In the given scenario, the Developers group needs the ability to create Azure Logic Apps within the Dev resource group. Assigning the Logic App Contributor role at the resource group level provides exactly that scope of control-users in the group can manage Logic App resources within that resource group without affecting other resources or permissions at the subscription level.
This satisfies the requirement under the principle of least privilege, as the Developers group gets only the permissions necessary to create and manage Logic Apps within Dev, and no more.
This is verified in the Microsoft Learn AZ-104 study guide under the topic "Manage Azure Role-Based Access Control (RBAC)", which emphasizes assigning the minimal level of role permissions required to complete administrative tasks at the appropriate scope (subscription, resource group, or resource).
Therefore, the proposed solution meets the goal because the Logic App Contributor role grants precisely the permissions needed.

NEW QUESTION: 197
You have an Azure virtual machine named VM1. VM1 was deployed by using a custom Azure Resource Manager template named ARM1.json.
You receive a notification that VM1 will be affected by maintenance.
You need to move VM1 to a different host immediately.

Solution: From the resource group blade, move VM1 to another resource group.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
Moving a virtual machine (VM) to a different host immediately in Azure is related to host-level maintenance mitigation, not to resource organization.
In this scenario, VM1 is affected by planned maintenance, and the goal is to move VM1 to a different host immediately.
According to Microsoft Azure Administrator documentation:
* Moving a VM to another resource group is an Azure Resource Manager (ARM) operation used for administrative organization.
* A resource group move does not change:
* The physical host
* The datacenter
* The availability zone
* The underlying hardware
* Therefore, moving VM1 to another resource group does not relocate the VM to a different host and does not avoid maintenance.
To move a VM to a different host in response to maintenance, supported options include:
* Redeploying the VM (which moves it to a new node)
* Using Availability Sets or Availability Zones
* Using Azure Scheduled Events to prepare workloads
Microsoft documentation explicitly states:
"Moving resources between resource groups does not change the location, availability zone, or the physical host of the resource." Because the proposed solution does not move the VM to a new host, it does not meet the goal.

NEW QUESTION: 198
You need to configure Azure Backup to meet the technical requirements for cont1 and share1.
To what should you set the backup frequency for each resource? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

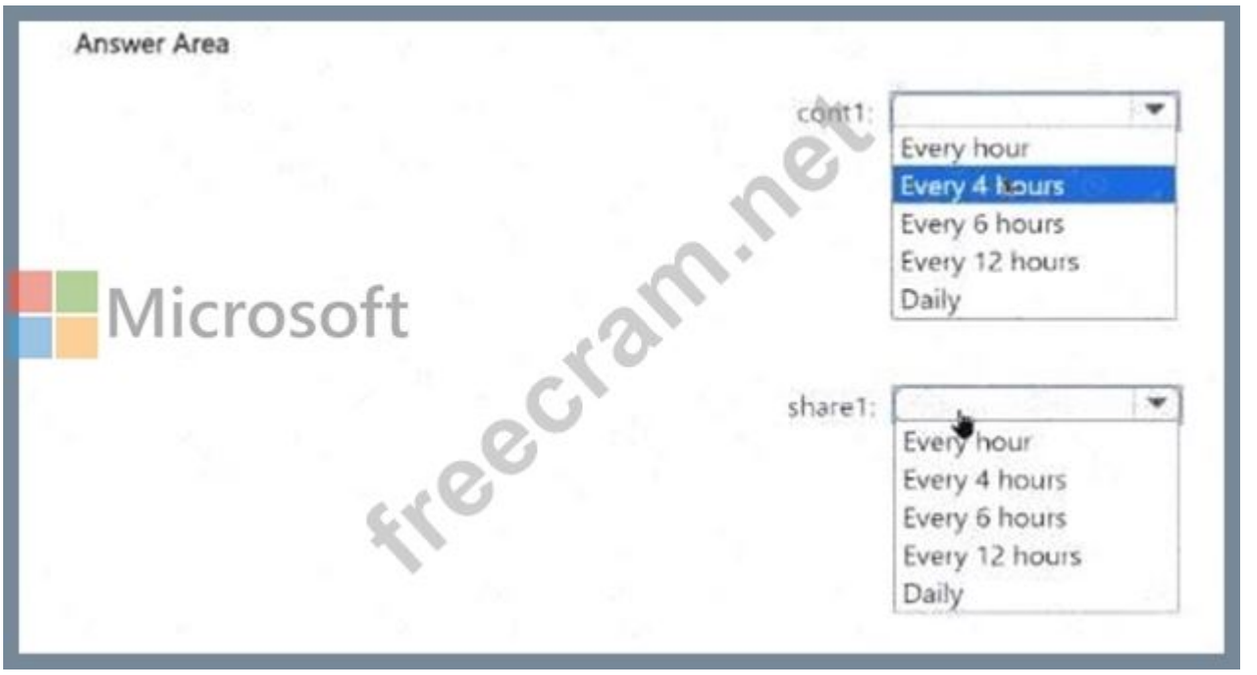

Answer:

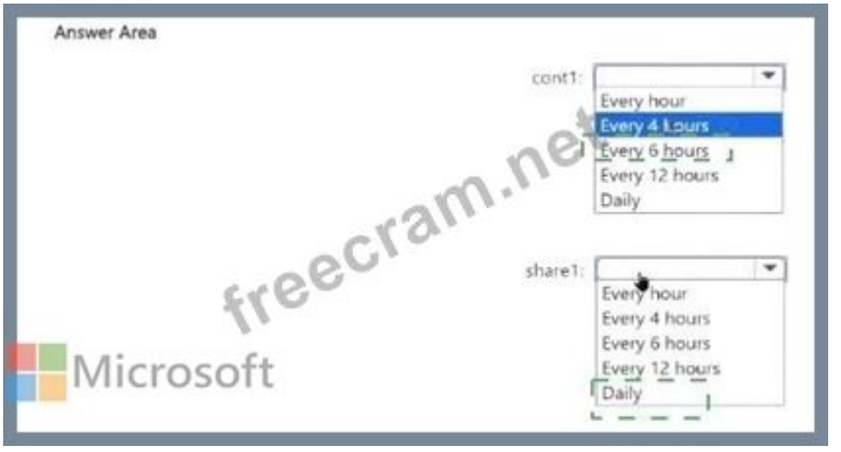

Explanation:

NEW QUESTION: 199

You have an Azure Linux virtual machine that is protected by Azure Backup.

One week ago, two files were deleted from the virtual machine.

You need to reses clients connect n on-premises computer as quickly as possible.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

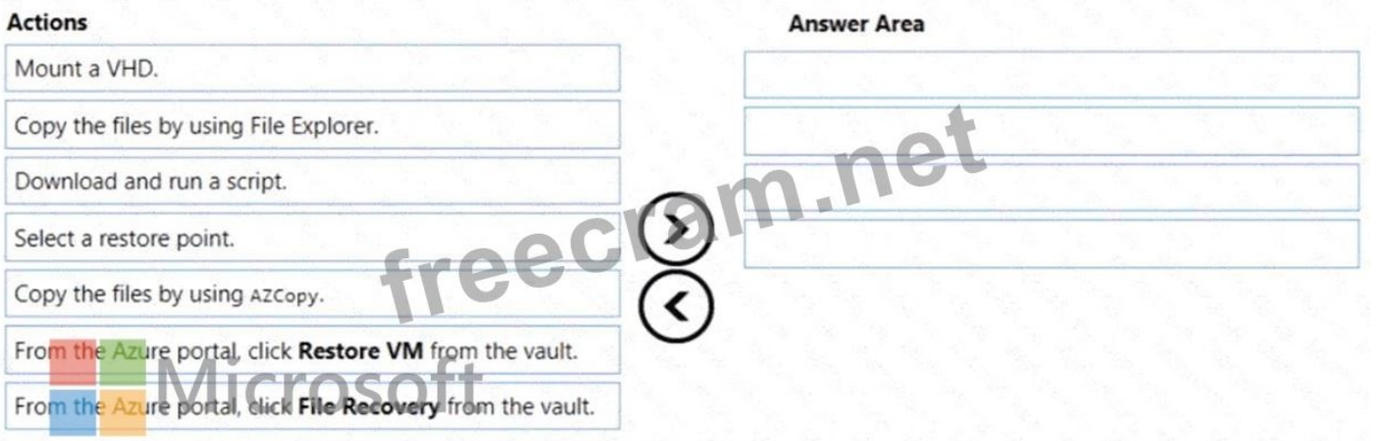

Answer:

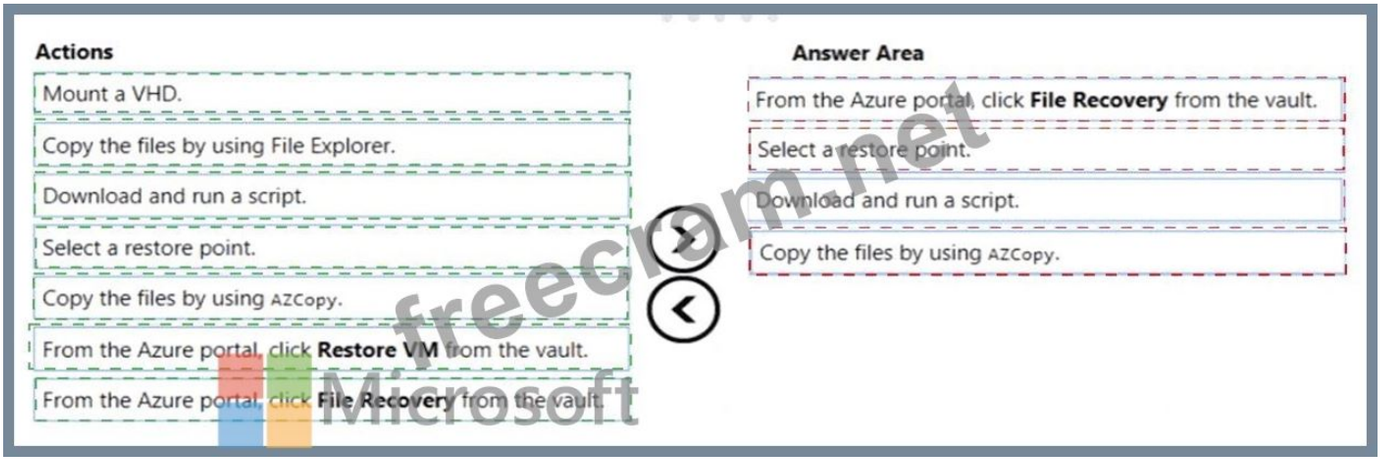

Explanation:

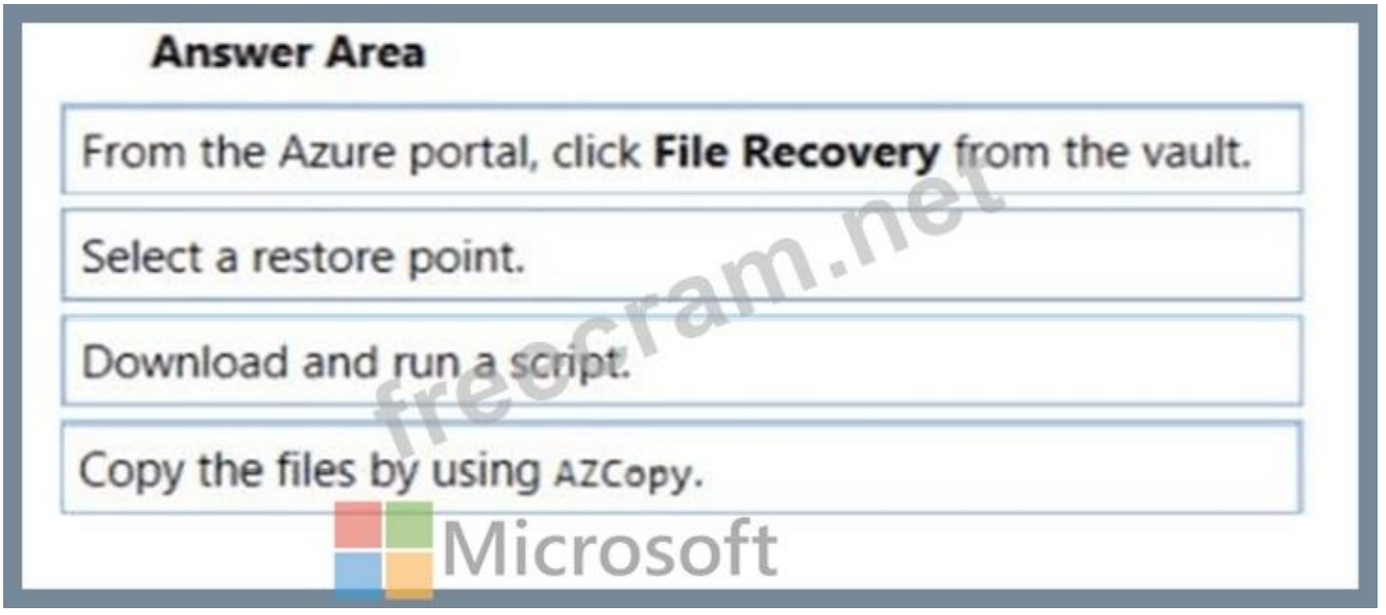

To restore files or folders from the recovery point, go to the virtual machine and choose the desired recovery point.

Step 0. In the virtual machine's menu, click Backup to open the Backup dashboard.

Step 1. In the Backup dashboard menu, click File Recovery.

Step 2. From the Select recovery point drop-down menu, select the recovery point that holds the files you want. By default, the latest recovery point is already selected.

Step 3: To download the software used to copy files from the recovery point, click Download

Executable (for Windows Azure VM) or Download Script (for Linux Azure VM, a python script is generated).
Step 4: Copy the files by using AzCopy
AzCopy is a command-line utility designed for copying data to/from Microsoft Azure Blob, File, and Table storage, using simple commands designed for optimal performance. You can copy data between a file system and a storage account, or between storage accounts.
References:
https://docs.microsoft.com/en-us/azure/backup/backup-azure-restore-files-from-vm
https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy