Microsoft.AZ-104.v2026-06-08.q181

Exam Code:	AZ-104
Exam Name:	Microsoft Azure Administrator Exam
Certification Provider:	Microsoft
Free Question Number:	181
Version:	v2026-06-08
# of views:	105
# of Questions views:	1815
https://www.freecram.net/torrent/Microsoft.AZ-104.v2026-06-08.q181.html	

NEW QUESTION: 1
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: 2
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: 3
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: 4
You have an Azure subscription that contains a virtual network named VNET1 in the East US 2 region.
Network Interfaces named VM1-NI and VM2-NI are connected to VNET1.

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

      "computerName": "VM2",
      "adminUsername": "AzureAdmin",
      "adminPassword": "[parameters('adminPassword')]"
    },
    "storageProfile": {
      "imageReference": "[variables('image')]",
      "osDisk": {
        "createOption": "FromImage"
      }
    },
    "networkProfile": {
      "networkInterfaces": [
```

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 and VM2 can connect to VNET1.	○	○
If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.	○	○
If the East US 2 region becomes unavailable, VM1 or VM2 will be available.	○	○

Answer:

Explanation:

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

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

NEW QUESTION: 5
You implement the planned changes for NSG1 and NSG2.
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

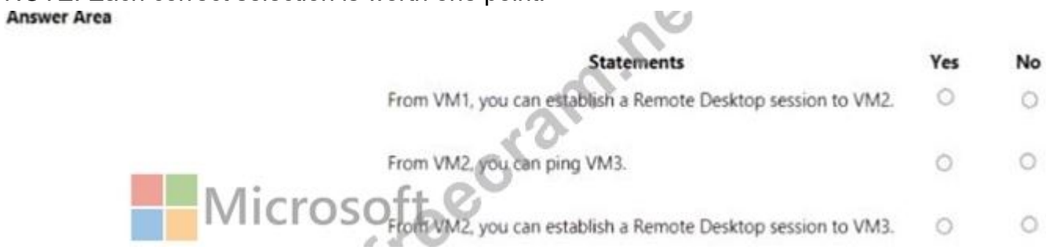

Answer:

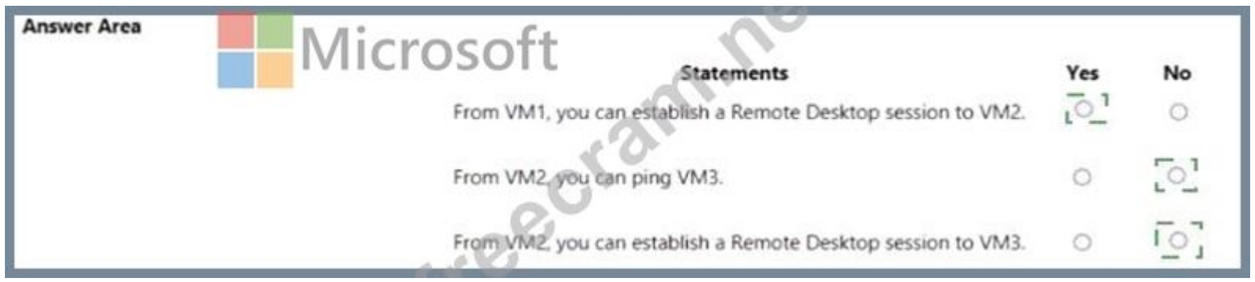

Explanation:

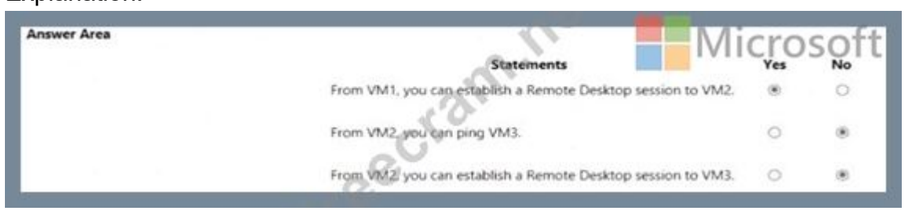

NEW QUESTION: 6
You have two Azure virtual networks named VNet1 and VNet2. VNet1 contains an Azure virtual machine named VM1. VNet2 contains an Azure virtual machine named VM2.
VM1 hosts a frontend application that connects to VM2 to retrieve data.
Users report that the frontend application is slower 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: 7

You have an Azure virtual machine named VM1 that connects to a virtual network named VNet1. VM1 has the following configurations:

Subnet: 10.0.0.0/24

Availability set: AVSet

Network security group (NSG): None

Private IP address: 10.0.0.4 (dynamic)

Public IP address: 40.90.219.6 (dynamic)

You deploy a standard, Internet-facing load balancer named slb1.

You need to configure slb1 to allow connectivity to VM1.

Which changes should you apply to VM1 as you configure slb1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

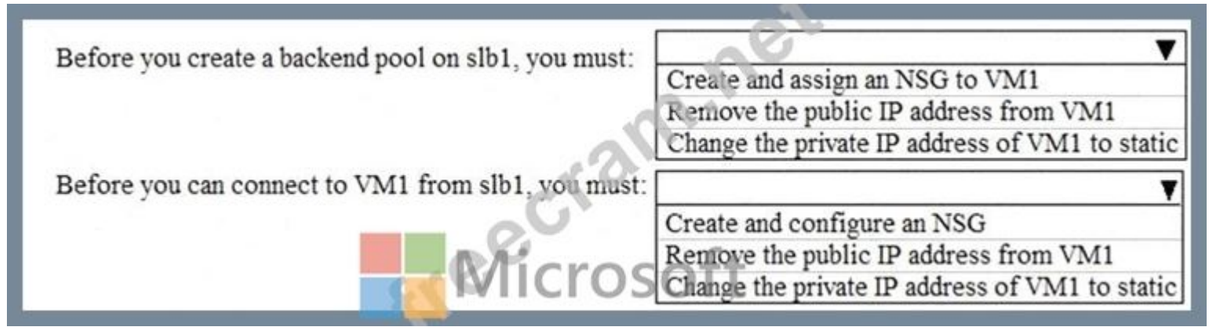

Answer:

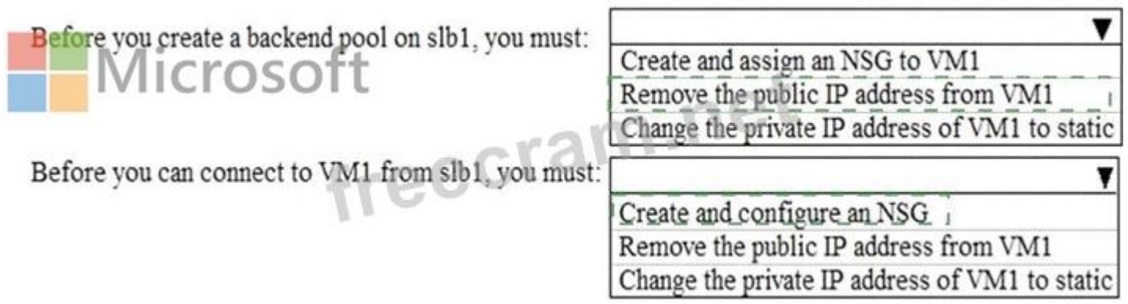

Explanation:

Box 1: Remove the public IP address from VM1

If the Public IP on VM1 is set to Dynamic, that means it is a Public IP with Basic SKU because Public IPs with Standard SKU have Static assignments by default, that cannot be changed. We cannot associate Basic SKUs IPs with Standard SKUs LBs. One cannot create a backend SLB pool if the VM to be associated has a Public IP. For Private IP it doesn ' t matter weather it is dynamic or static, still we can add the such VM into the SLB backend pool.

Box 2: Create and configure an NSG

Standard Load Balancer is built on the zero trust network security model at its core. Standard Load Balancer secure by default and is part of your virtual network. The virtual network is a private and isolated network.

This means Standard Load Balancers and Standard Public IP addresses are closed to inbound flows unless opened by Network Security Groups. NSGs are used to explicitly permit allowed traffic. If you do not have an NSG on a subnet or NIC of your virtual machine resource, traffic is not allowed to reach this resource. To learn more about NSGs and how to apply them for your scenario, see Network Security Groups. Basic Load Balancer is open to the internet by default.

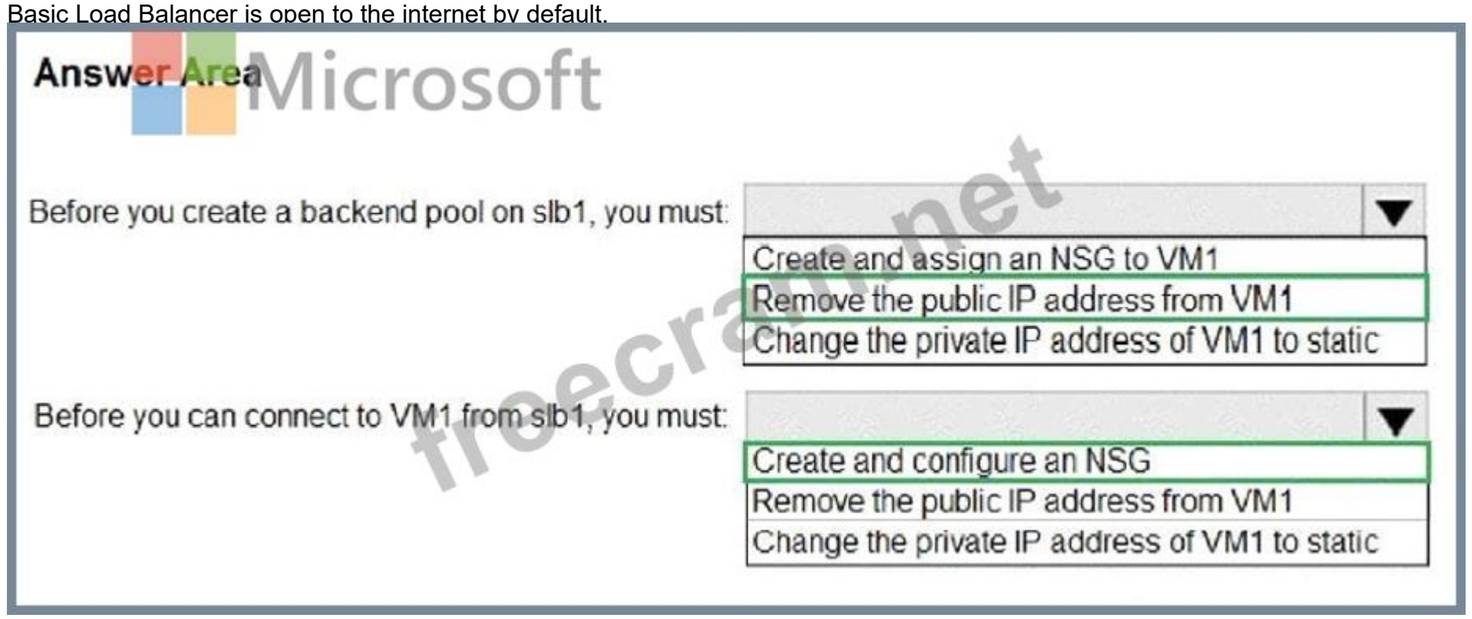

Reference:
https://docs.microsoft.com/en-us/azure/load-balancer/quickstart-load-balancer-standard-public-portal
https://docs.microsoft.com/en-us/azure/load-balancer/load-balancer-overview

NEW QUESTION: 8
You have an Azure subscription that contains the resource groups shown in the following table.
RG1 contains the resources shown in the following table.

Name	Type	Lock name	Lock type
storage1	Storage account	Lock1	Delete
VNET1	Virtual network	Lock2	Read-only
IP1	Public IP address	*None*	*None*

RG2 contains the resources shown in the following table.

Name	Type	Lock name	Lock type
storage2	Storage account	Lock1	Delete
VNET2	Virtual network	Lock2	Read-only
IP2	Public IP address	*None*	*None*

You need to identify which resources you can move from RG1 to RG2, and which resources you can move from RG2 to RG1. Which resources should you identify? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

NEW QUESTION: 9
You have an Azure subscription.
You plan to create an Azure container registry named ContReg1.
You need to ensure that you can push and pull signed images for ContReg1. What should you do for ContReg1?
A. Enable content trust.
B. Add a token.
C. Enable encryption by using a customer-managed key.
D. Create a connected registry.
Answer: (SHOW ANSWER)
To push and pull signed container images in Azure Container Registry (ACR), you must enable content trust.
Content trust integrates with Docker Notary and allows image publishers to sign images and consumers to verify those signatures during pull operations.
The other options do not meet the requirement:
* Add a token # controls authentication, not image signing
* Enable encryption with customer-managed keys # protects data at rest, not image signatures
* Create a connected registry # supports edge/replication scenarios, not signing

NEW QUESTION: 10
You have an Azure subscription.
You create the following file named Deploy.json.

```
        "sku": {
          "name": "Premium_LRS"
        },
        "kind": "StorageV2",
        "properties": {},
        "copy": {
          "name": "storagecopy",
          "count": 3
        }
      }

  ]
}
```

You connect to the subscription and run the following commands.

```
New-AzResourceGroup -Name RG1 -Location "centralus"
New-AzResourceGroupDeployment -ResourceGroupName RG1 -TemplateFile "deploy.json"
```

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
The commands will create four new resources.	○	○
The commands will create storage accounts in the West US Azure region.	○	○
The first storage account that is created will have a prefix of 0.	○	○

Answer:

Answer Area

Statements	Yes	No
The commands will create four new resources.	○	○
The commands will create storage accounts in the West US Azure region.	○	○
The first storage account that is created will have a prefix of 0.	○	○

Explanation:

Answer Area

Statements	Yes	No
The commands will create four new resources.	●	○
The commands will create storage accounts in the West US Azure region.	○	●
The first storage account that is created will have a prefix of 0.	●	○

NEW QUESTION: 11

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.

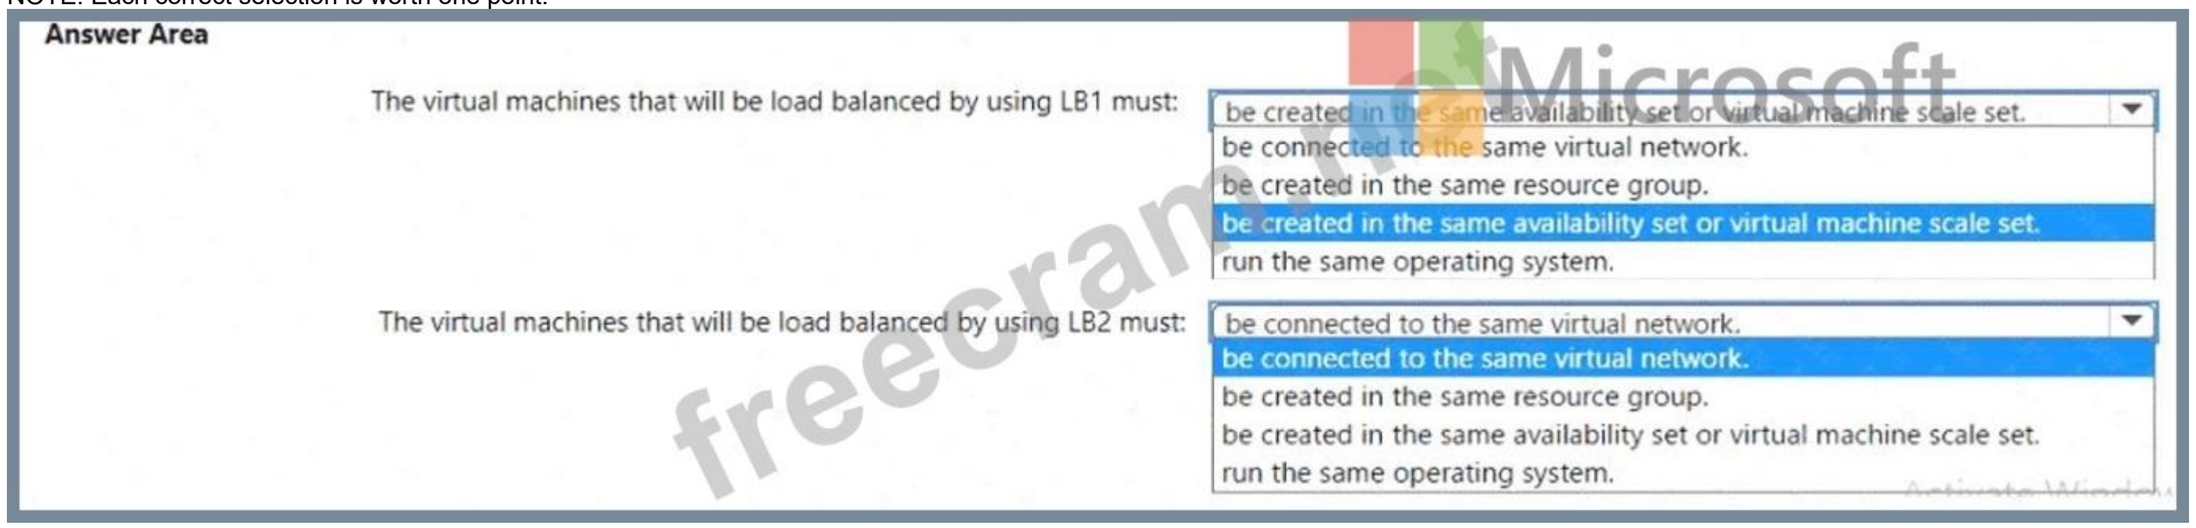

Answer:

Answer Area

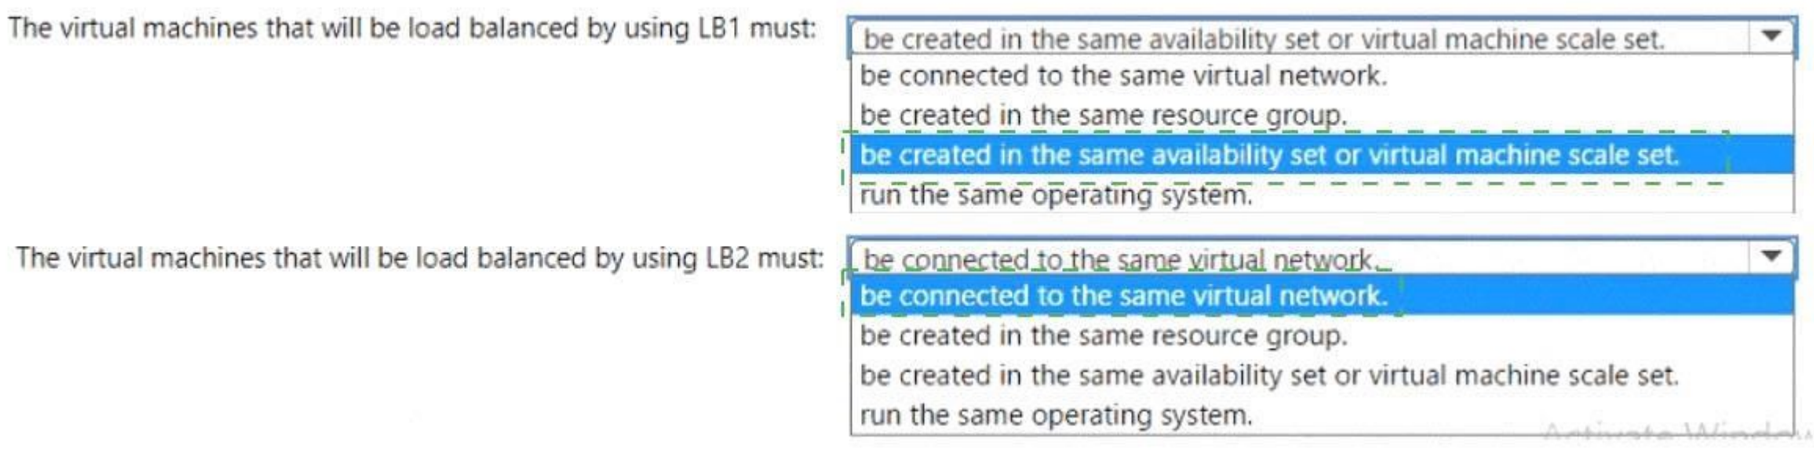

Explanation:

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: 12

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: 13
You have an Azure Storage accounts as 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 Area

You can use [**answer choice**] for Azure Table Storage.

- storageaccount1 only
- storageaccount2 only
- storageaccount3 only
- storageaccount1 and storageaccount2 only
- storageaccount2 and storageaccount3 only

You can use [**answer choice**] for Azure Blob storage.

- storageaccount3 only
- storageaccount2 and storageaccount3 only
- storageaccount1 and storageaccount3 only
- all the storage accounts

Answer:

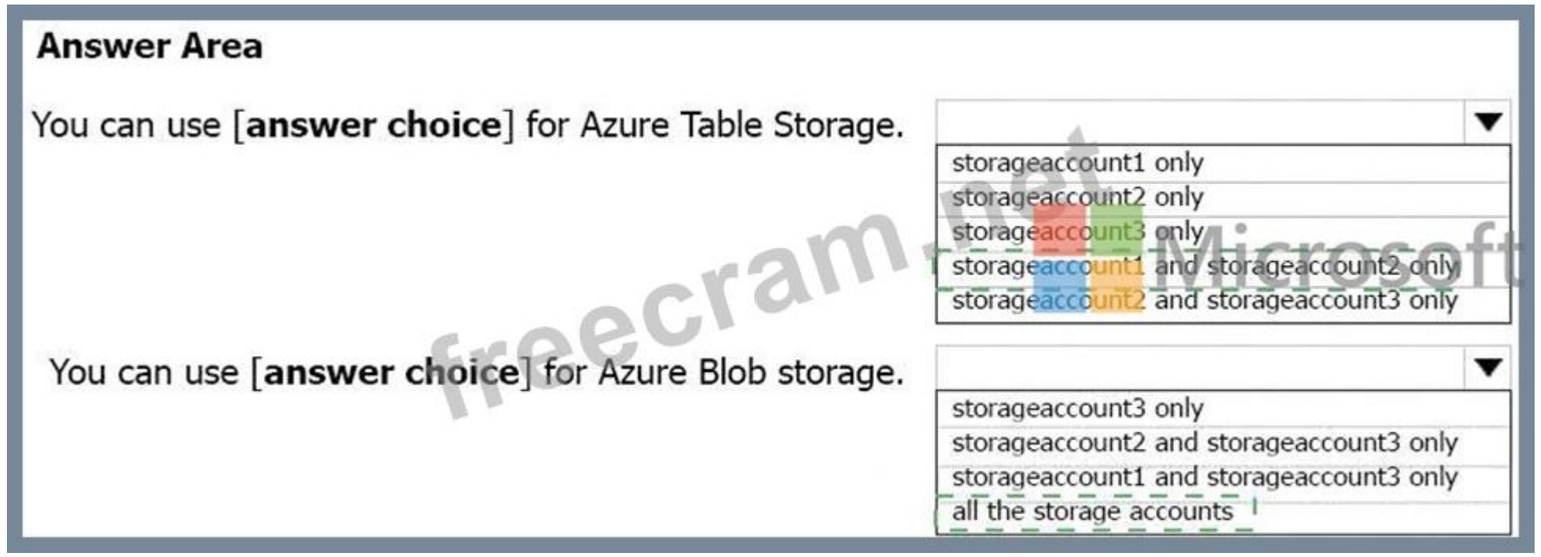

Explanation:

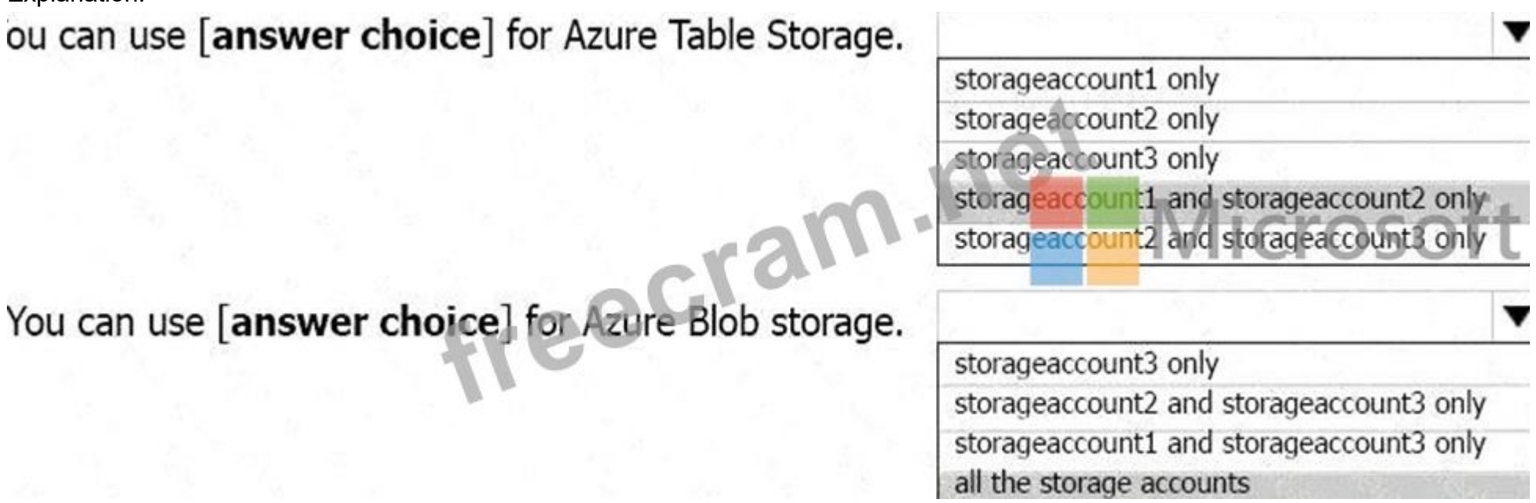

Box 1: storageaccount1 and storageaccount2 only

Box 2: All the storage accounts

Note: The three different storage account options are: General-purpose v2 (GPv2) accounts, General-purpose v1 (GPv1) accounts, and Blob storage accounts.

General-purpose v2 (GPv2) accounts are storage accounts that support all of the latest features for blobs, files, queues, and tables.

Blob storage accounts support all the same block blob features as GPv2, but are limited to supporting only block blobs.

General-purpose v1 (GPv1) accounts provide access to all Azure Storage services, but may not have the latest features or the lowest per gigabyte pricing.

References: https://docs.microsoft.com/en-us/azure/storage/common/storage-account-options

NEW QUESTION: 14

You have an Azure Storage account named storage1.

You have Azure App Service apps named App1 and App2 that run in an Azure container instance. Each app uses a managed identity.

You need to ensure that App1 and App2 can read blobs from storage1. The solution must meet the following requirements:

* Minimize the number of secrets used.
* Ensure that App2 can only read from storage1 for the next 30 days.

What should you configure in storage1 for each app? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point

Answer Area

Answer:

Answer Area

Explanation:

Answer Area

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

Let's analyze the requirements and correct configuration for each app based on Azure's security and access control models.
App1 - Minimize Secrets
App1 uses a managed identity, meaning it can be authenticated to Azure services without any stored credentials or secrets.
The best practice is to assign Azure RBAC permissions (role-based access control) directly at the storage account or container level.
By using Access control (IAM), you can assign the Storage Blob Data Reader role to App1's managed identity.
This method uses Azure AD-based authentication, requires no SAS tokens or access keys, and minimizes secret management.
Access is continuous until the role is removed or modified.
Therefore, App1 # Access control (IAM)
App2 - Temporary 30-day Access
The requirement specifies that App2 should be able to read blobs only for 30 days.
Azure RBAC roles (IAM) do not provide time-bound permissions.
The appropriate way to grant time-limited access is through a Shared Access Signature (SAS).
A SAS token defines permissions, resource scope (e.g., container or blob), and an expiry time - making it ideal for temporary or limited access scenarios.
You can generate a SAS token valid for 30 days and assign it to App2.
Therefore, App2 # Shared access signatures (SAS)
Why Not Access Keys or Advanced Security
Access Keys: Grant full control (read/write/delete) to the storage account - not secure or granular, and they cannot be time-bound.
Advanced Security: Refers to configurations such as firewall rules or encryption; not directly related to granting app access.
Microsoft Azure Administrator Documentation Extract (AZ-104 Study Guide Reference):
"To enable secure access for applications, use Azure AD authentication with managed identities and assign appropriate RBAC roles via Access control (IAM). For temporary or limited access, use Shared Access Signatures (SAS) to specify permissions and expiry times." (Source: Microsoft Learn - Secure access to Azure Storage with Azure AD, SAS, and managed identities.)
Final Verified Answer:
App1: Access control (IAM)
App2: Shared access signatures (SAS)

NEW QUESTION: 15
Your network contains an on-premises Active Directory Domain Services (AD DS) domain named contoso.
com. The domain contains the servers shown in the following table.

Name	IP address	Role
DC1	192.168.2.1/16	Domain controller DNS server
Server1	192.168.2.50/16	Member server

You plan to migrate contoso.com to Azure.
You create an Azure virtual network named VNET1 that has the following settings:
* Address space: 10.0.0.0/16
* Subnet:
o Name: Subnet1
o IPv4: 10.0.1.0/24
You need to move DC1 to VNET1. The solution must ensure that the member servers in contoso.com can resolve AD DS DNS names.
How should you configure DC1? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Answer:

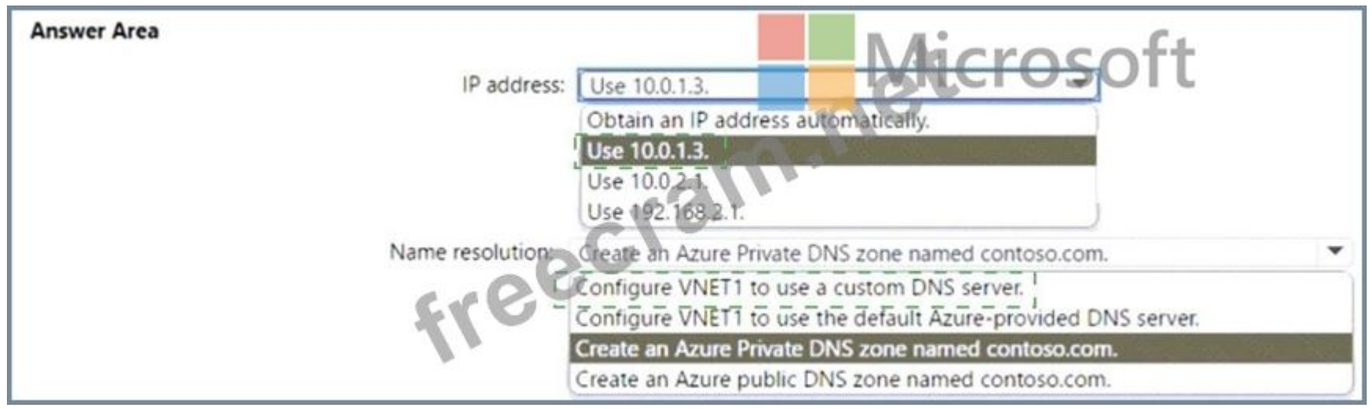

Explanation:

IP address: You should use 10.0.1.3 as the IP address for DC1. This is because DC1 needs to have a static IP address within the subnet range of VNET1, which is 10.0.1.0/241. You cannot use 10.0.2.1 or 192.168.2.1, as they are outside of the subnet range of VNET1. You also cannot obtain an IP address automatically, as this may cause DC1 to lose its IP address and break the DNS resolution for the domain members2.

Name Resolution: You should configure VNET1 to use a custom DNS server that points to the IP address of DC1, which is 10.0.1.33. This is because DC1 is the domain controller and DNS server for contoso.com, and it needs to resolve the AD DS DNS names for the domain members that are in Azure or on-premises. You cannot use the default Azure-provided DNS server, as it does not support AD DS DNS names. You also do not need to create an Azure Private DNS zone or an Azure public DNS zone named contoso.com, as these are not required for AD DS DNS resolution.

NEW QUESTION: 16

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: (SHOW ANSWER)

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: 17
You have an Azure subscription.
You need to use an Azure Resource Manager (ARM) template to create a virtual machine that will have multiple data disks.
How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

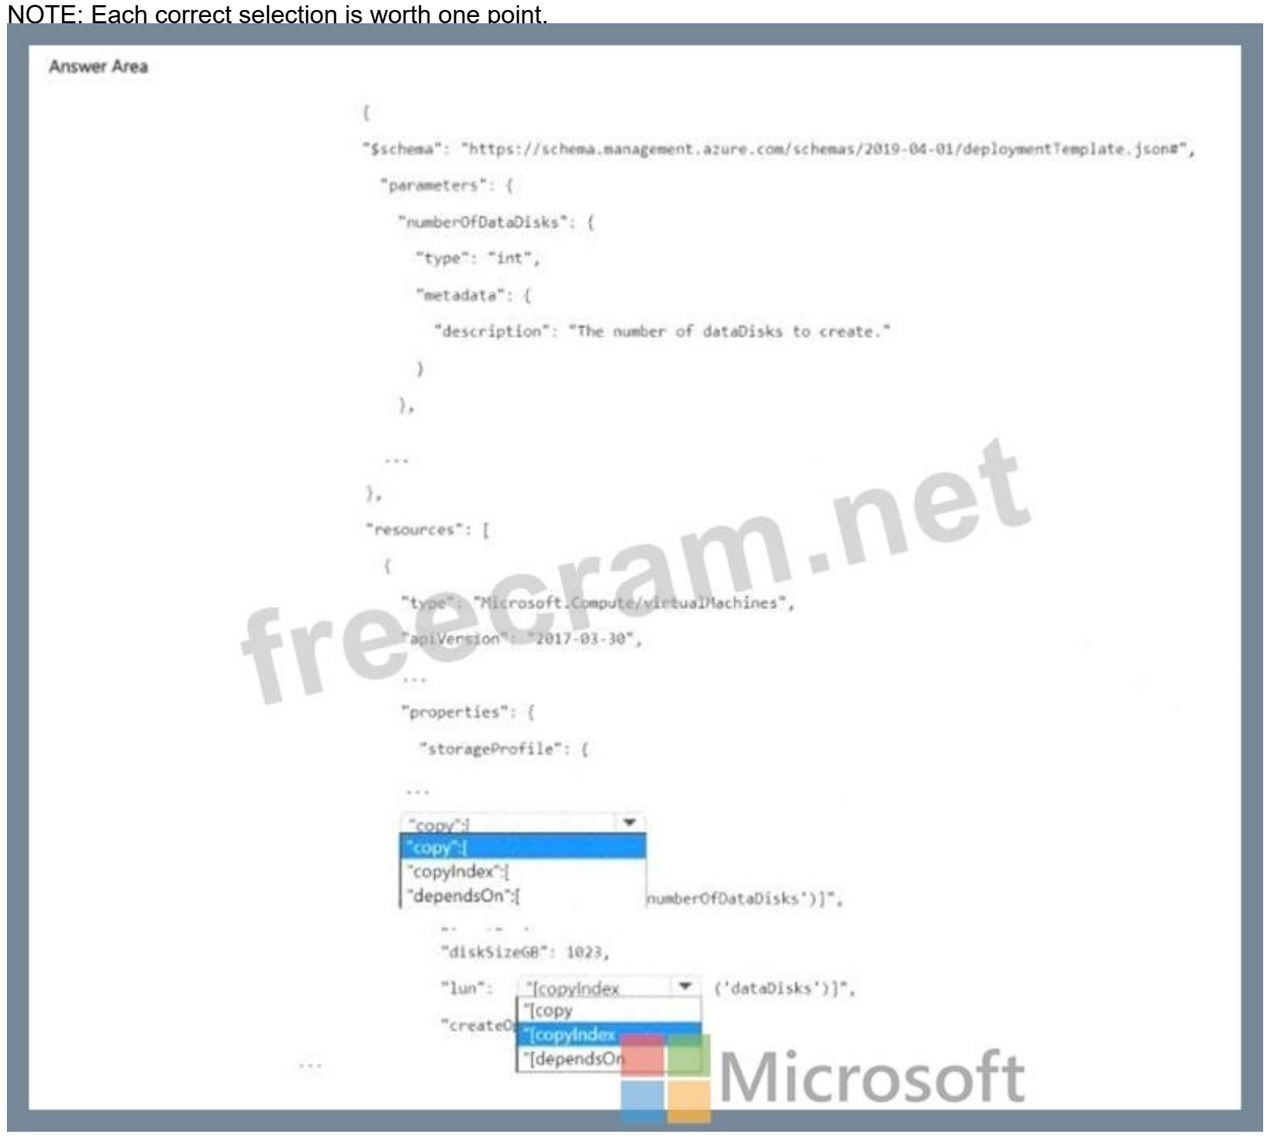

Answer:

Answer Area

Explanation:

Answer Area

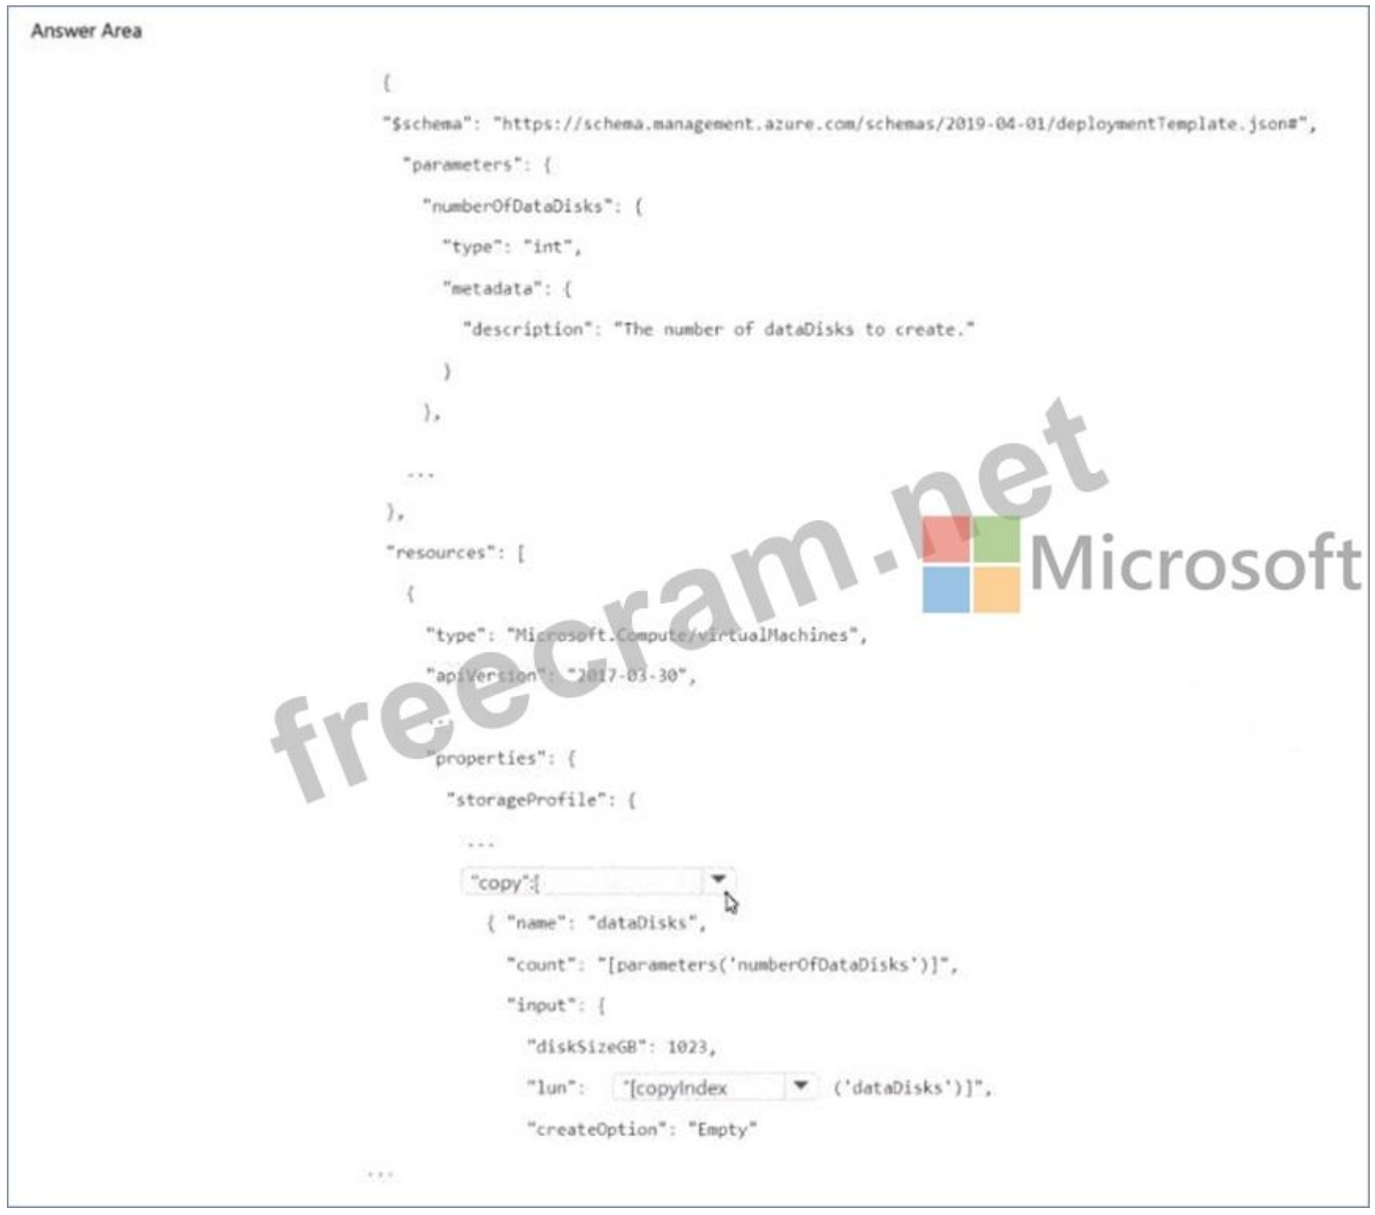

When using an Azure Resource Manager (ARM) template to deploy multiple identical resources-such as several data disks for a virtual machine-you use the copy loop construct within the resource definition.

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: 18
You have an Azure subscription named Subscription 1 and an on-premises deployment of Microsoft System Center Service Manager Subscription! contains a virtual machine named VM1.
You need to ensure that an alert is set in Service Manager when the amount of available memory on VM1 is below 10 percent. What should you do first?
A. Create a notification.
B. Create an automation runbook.
C. Deploy the IT Service Management Connector (ITSM).
D. Deploy a function app
Answer: (SHOW ANSWER)
IT Service Management Connector (ITSMC) allows you to connect Azure to a supported IT Service Management (ITSM) product or service. Azure services like Azure Log Analytics and Azure Monitor provide tools to detect, analyze, and troubleshoot problems with your Azure and non-Azure resources. But the work items related to an issue typically reside in an ITSM product or service. ITSMC provides a bi-directional connection between Azure and ITSM tools to help you resolve issues faster. ITSMC supports connections with the following ITSM tools: ServiceNow, System Center Service Manager, Provance, Cherwell.

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

NEW QUESTION: 19
You have an Azure subscription that contains a storage account named storage1. The subscription is linked to an Azure Active Directory (Azure AD) tenant named contoso.com that syncs to an on-premises Active Directory domain.
The domain contains the security principals shown in the following table.

Name	Type
User1	User
Computer1	Computer

In Azure AD, you create a user named User2.
The storage1 account contains a file share named share1 and has the following configurations.

```
"kind": "StorageV2",
"properties": {
                "azureFilesIdentityBasedAuthentication": {
                    "directoryServiceOptions": "AD",
                    "activeDirectoryProperties": {
                        "domainName": "Contoso.com",
                        "netBiosDomainName": "Contoso.com",
                        "forestName": "Contoso.com",
-
                    }
            }
```

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
You can assign the Storage File Data SMB Share Contributor role to User1 for share1.	○	○
You can assign the Storage File Data SMB Share Reader role to Computer1 for share1.	○	○
You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1.	○	○

Answer:

Statements	Yes	No
You can assign the Storage File Data SMB Share Contributor role to User1 for share1.	○	○
You can assign the Storage File Data SMB Share Reader role to Computer1 for share1.	○	○
You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1.	○	○

Explanation:

Statements	Yes	No
You can assign the Storage File Data SMB Share Contributor role to User1 for share1.	○	○
You can assign the Storage File Data SMB Share Reader role to Computer1 for share1.	○	○
You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1.	○	○

Reference:
https://docs.microsoft.com/en-us/azure/storage/files/storage-files-identity-ad-ds-assign-permissions?tabs=azure-portal

Azure Files supports Azure Active Directory Domain Services (Azure AD DS) authentication for SMB access, allowing you to use your existing Active Directory (AD) identities and groups to control access to Azure file shares.In this scenario, the configuration snippet shows that the storage account (storage1) is using Active Directory (AD) authentication through azureFilesIdentityBasedAuthentication with directoryServiceOptions: " AD " .

According to Microsoft Azure documentation, the Storage File Data SMB Share built-in roles (Reader, Contributor, Elevated Contributor) can only be assigned to Azure AD users, groups, or service principals, not to computer objects.

Statement 1: # User1 is an on-premises AD user synchronized to Azure AD. Since Azure Files uses identity- based authentication through AD, User1 can be assigned roles like Storage File Data SMB Share Contributor to access share1.

Statement 2: # Computer1 is a computer account and cannot have Azure RBAC roles directly assigned in Azure. Therefore, you cannot assign SMB Share Reader to it.

Statement 3: # User2 is an Azure AD user (cloud-only identity). Azure RBAC roles for storage shares can also be assigned to Azure AD users. Hence, you can assign Storage File Data SMB Share Elevated Contributor to User2.

NEW QUESTION: 20

You create an Azure Storage account named Contoso storage.
You plan to create a file share named data.
Users need to map a drive to the data file share from home computers that run Windows 10.
Which outbound port should be open between the home computers and the data file share?

A. 80
B. 443
C. 445
D. 3389

Answer: (SHOW ANSWER)

When users map an Azure file share (hosted in Azure Files) to their Windows 10 devices using SMB (Server Message Block) protocol, the connection uses TCP port 445 for communication.

According to the official Microsoft Azure Files documentation ("Port requirements for Azure file shares" and
"Use an Azure file share with Windows"), the SMB protocol version 3.0 is used for secure, encrypted communication between the client and the Azure Files service. SMB communicates over TCP port 445, which must be open outbound from the client system to the Azure Storage endpoint (for example, contosostorage.
file.core.windows.net).

Here's how Azure handles access:

* SMB over TCP port 445 is used for file shares (for mounting and drive mapping).
* HTTPS (port 443) is used for REST API and portal access to the storage account.
* Port 80 is optional (used only for unencrypted HTTP requests, not supported for SMB access).
* Port 3389 is for Remote Desktop Protocol (RDP), unrelated to Azure Files.

Therefore, to allow users to map the Azure Files share (\contosostorage.file.core.windows.net\data) from their home computers, TCP port 445 must be open on outbound connections.

NEW QUESTION: 21

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

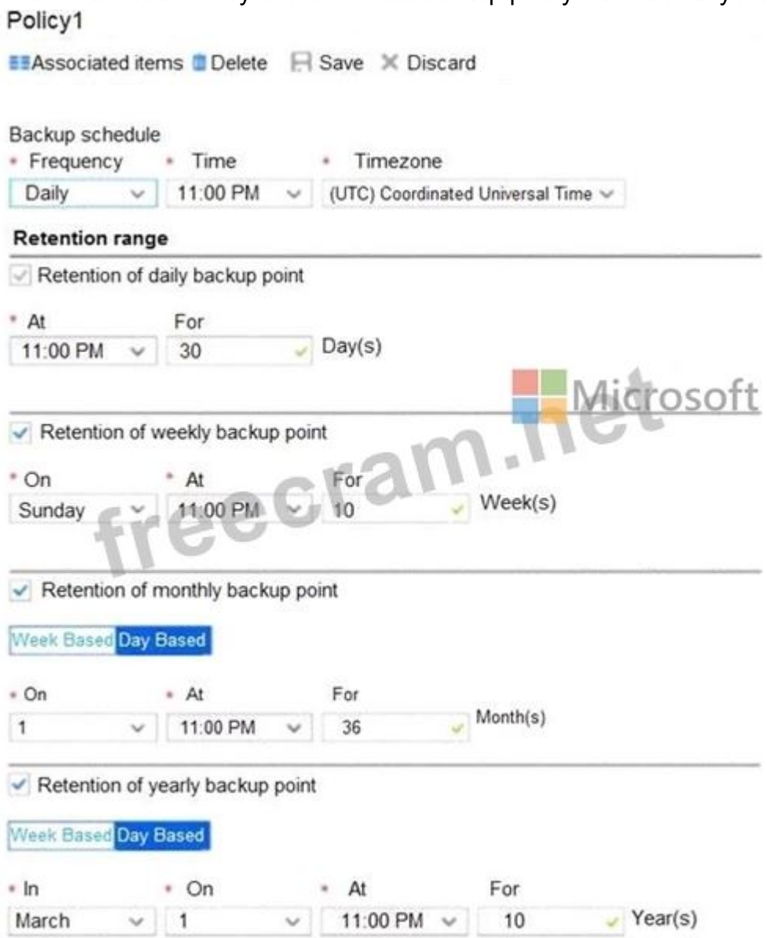

Answer Area

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 Area

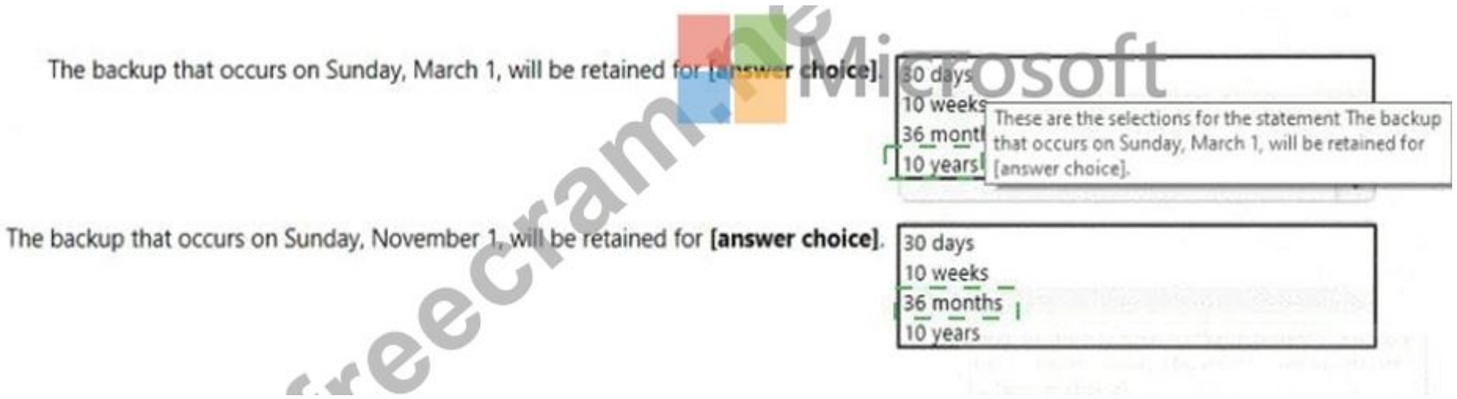

Explanation:
Box 1: 10 years
The yearly backup point occurs to 1 March and its retention period is 10 years.
Box 2: 36 months
The monthly backup point occurs on the 1
of every month and its retention period is 36 months.
Note: Azure retention policy takes the longest period of retention for each backup.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

NEW QUESTION: 22
You have an Azure subscription.
You plan to use Azure Resource Manager templates to deploy 50 Azure virtual machines that will be part of the same availability set.
You need to ensure that as many virtual machines as possible are available if the fabric fails or during servicing.
How should you configure the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

```
"$schema": https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json
 "contentVersion" : "1.0.0.0",
 "parameters": {},
 "resources": [
   {
        "type": "Microsoft.Compute/availabilitySets",
        "name": "ha",
        "apiVersion": "2017-12-01",
        "location": "eastus",
        "properties": {
            "platformFaullDomainCount":  ,
                                          max value
                                          0
                                          20
           "platformUpdateDomainCount":
           }                              max value
         }                                0
      ]                                   20
   }
```


Answer:

```
{
  "$schema": https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json,
   "contentVersion" : "1.0.0.0",
  "parameters": {},
  "resources": [
    {
         "type": "Microsoft.Compute/availabilitySets",
         "name": "ha",
         "apiVersion": "2017-12-01",
         "location": "eastus",
         "properties": {
             "platformFaullDomainCount":  ,
                                          max value
                                          0
                                          20
            "platformUpdateDomainCount":
            }                             max value
          }                               0
       ]                                  20
    }
```


Explanation:

Box 1 = max value

Box 2 = 20

Explanation:
Use max for platformFaultDomainCount
2 or 3 is max value, depending on which region you are in.
Use 20 for platformUpdateDomainCount
Increasing the update domain (platformUpdateDomainCount) helps with capacity and availability planning when the platform reboots nodes. A higher number for the pool (20 is max) means that fewer of their nodes in any given availability set would be rebooted at once.
References:
https://www.itprotoday.com/microsoft-azure/check-if-azure-region-supports-2-or-3-fault-domains-managed- disks
https://github.com/Azure/acs-engine/issues/1030

NEW QUESTION: 23
You have an Azure subscription that contains a storage account named storage.
You have the devices shown in the following table.

Name	Platform
Device1	Windows 10
Device2	Linux
Device3	macOS

From which devices can you use AzCopy to copy data to storage1?
A. Device1 and Device2 only
B. Device1, Device2 and Device3
C. Device' only
D. Device and Device3 only
Answer: (SHOW ANSWER)
AzCopy is a Microsoft command-line utility designed for copying data to and from Azure Storage accounts, including Blob Storage, File Storage, and Table Storage. It is part of the Azure Storage Tools suite and is widely used by administrators to perform high-performance data migration and synchronization operations.
According to the Microsoft Azure Administrator study guide and official Azure documentation, AzCopy is a cross-platform tool. The latest versions (from AzCopy v10 onwards) are developed using Go (Golang), which enables native support for multiple operating systems. Specifically, Microsoft provides official binaries for Windows, Linux, and macOS operating systems.
Here is a detailed extract based on Microsoft's official documentation:
"AzCopy is a command-line utility that you can use to copy blobs or files to or from a storage account.
AzCopy is available on Windows, Linux, and macOS platforms. You can download the tool and install it manually or use it via Azure Cloud Shell where it's preinstalled." This means that Device1 (Windows 10), Device2 (Linux), and Device3 (macOS) can all use AzCopy to copy data to or from Azure Storage.
There are no licensing or configuration restrictions that limit its operation to any specific OS, and all three platforms support authentication methods such as Azure AD credentials, SAS tokens, and storage account keys.
Hence, the verified and Microsoft-documented correct answer is:

NEW QUESTION: 24
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: 25

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 Area

Answer:

Answer Area

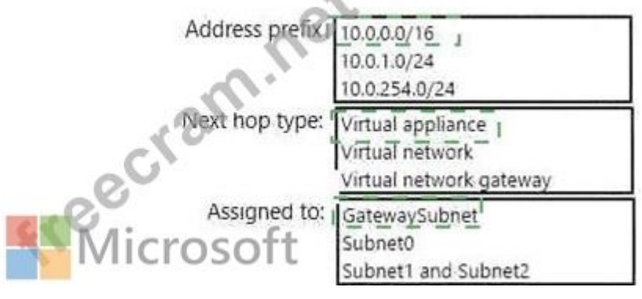

Explanation:

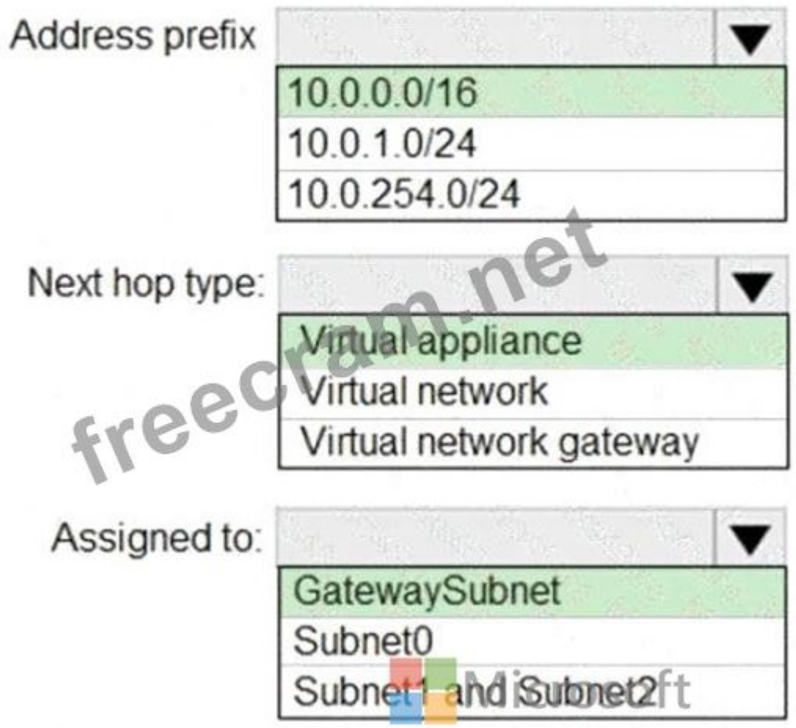

Box1 : 10.0.0.0/16

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: 26

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:

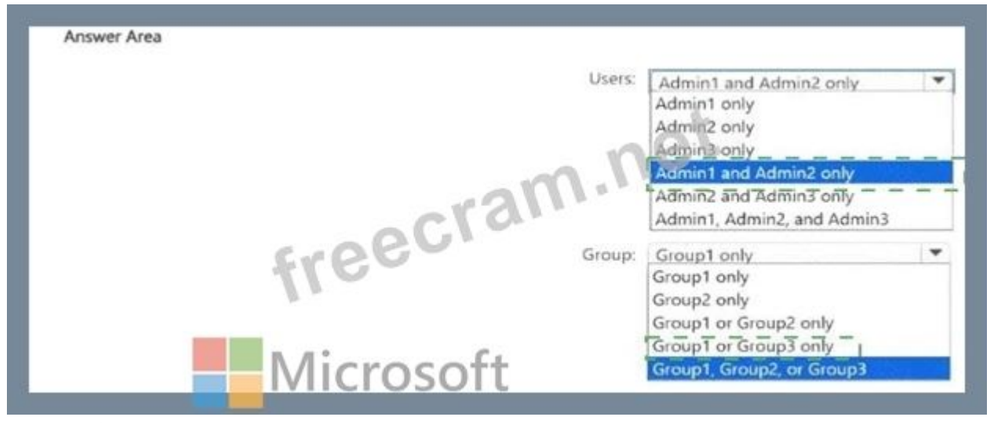

Explanation:

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: 27
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: 28
You have an Azure subscription that has the public IP addresses shown in the following table.

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

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

NEW QUESTION: 29
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: 30
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. Increase the resource limit of container2 to six CPUs.
B. Decrease the resource limit of container2 to two CPUs.
C. Increase the resource limit of container1 to three CPUs.
D. Remove the resource limit for both containers.
Answer: (SHOW ANSWER)

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 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: (SHOW ANSWER)
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: 32
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have 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: 33
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.)

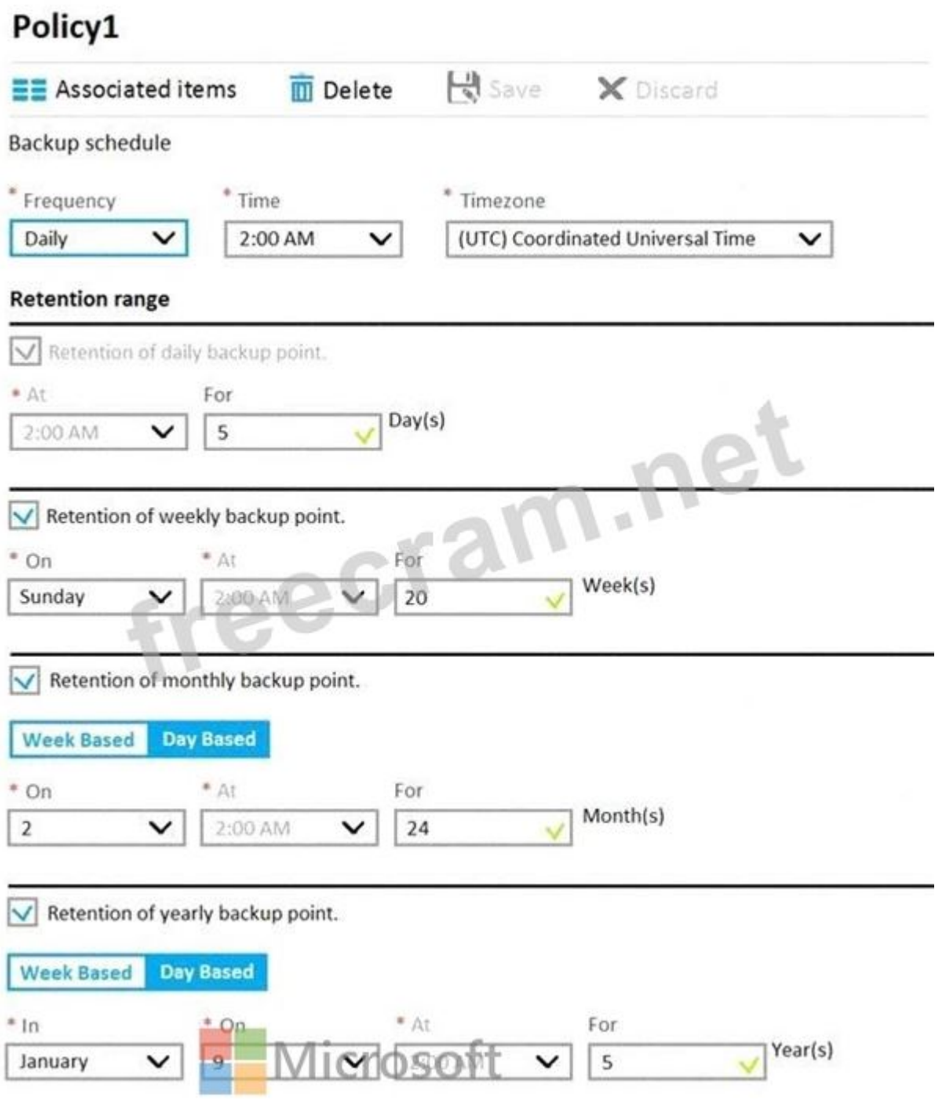

You configure the backup of VM1 to use Policy1 on Thursday, January 1.

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.

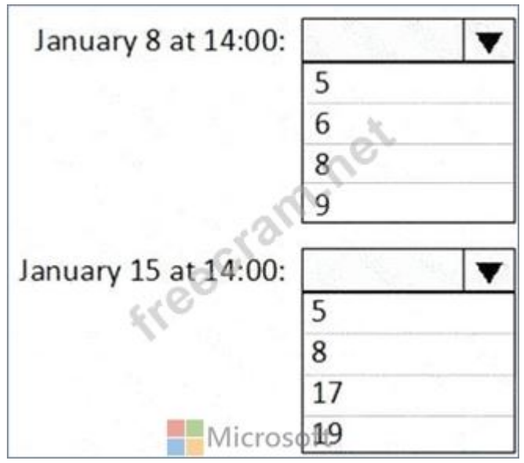

Answer:

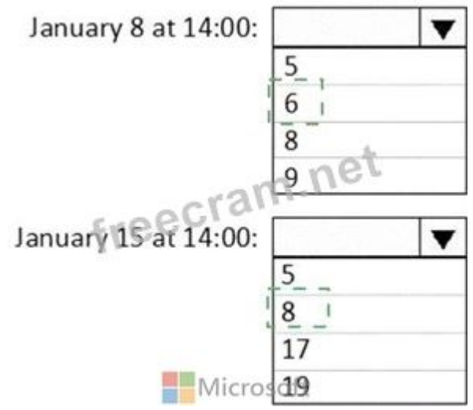

Explanation:

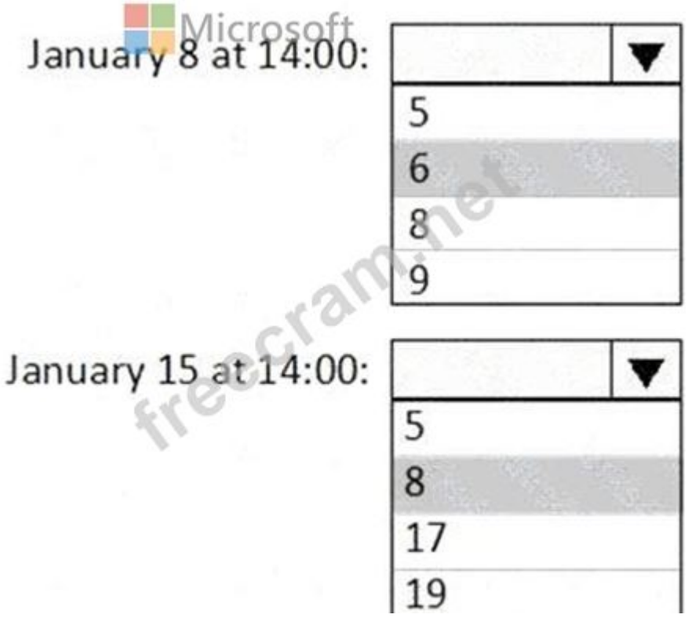

NEW QUESTION: 34
You plan to automate the deployment of a virtual machine scale set that uses the Windows Server 2016 Datacenter image.
You need to ensure that when the scale set virtual machines are provisioned, they have web server components installed.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE Each correct selection is worth one point.
A. Modify the extensionProfile section of the Azure Resource Manager template.
B. Create a new virtual machine scale set in the Azure portal.
C. Create an Azure policy.
D. Create an automation account.
E. Upload a configuration script.
Answer: (SHOW ANSWER)
To automate the deployment of a virtual machine scale set that uses the Windows Server 2016 Datacenter image and has web server components installed, you need to perform the following actions:
Modify the extensionProfile section of the Azure Resource Manager template. This section defines the extensions that are applied to the scale set virtual machines after they are provisioned. You can use the Custom Script Extension to run PowerShell scripts that install and configure the web server components. For more information, see Deploy an application to an Azure Virtual Machine Scale Set1.
Upload a configuration script. This is the PowerShell script that contains the commands to install and configure the web server components. You can upload the script to a storage account or a GitHub repository, and then reference it in the extensionProfile section of the template. For an example of a configuration script, see Tutorial: Install applications in Virtual Machine Scale Sets with Azure PowerShell2.

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.	○	○

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 plan to deploy an Azure container instance by using the following Azure Resource Manager template.

```
{
    "type": "Microsoft.ContainerInstance/containerGroups",
    "apiVersion": "2018-10-01",
    "name": "webprod",
    "location": "westus",
    "properties": {
        "containers": [
            {
                "name": "webprod",
                "properties": {
  "image": "microsoft/iis:nanoserver",
  "ports": [
     {
         "protocol": "TCP",
         "port": 80
     }

          "environmentVariables": [],
          "resources": {
              "requests": {
                  "memoryInGB": 1.5,
                  "cpu": 1
              }
          }
      }
    }
],
"restartPolicy": "OnFailure",
"ipAddress": {
    "ports": [
                    }
                ],
                "ip": "[parameters('IPAddress')]",
                "type": "Public"
            },
            "osType": "Windows"
        }
    }
```

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

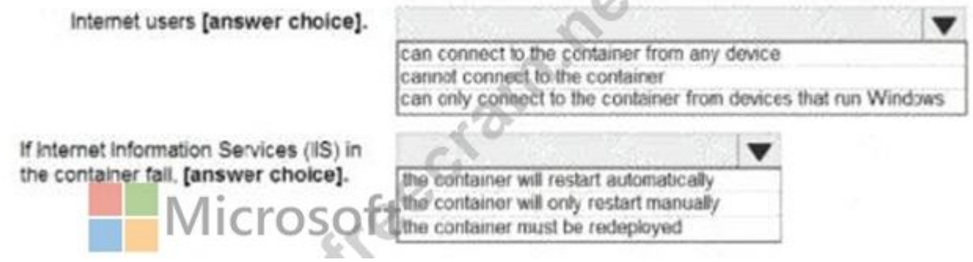

Answer:

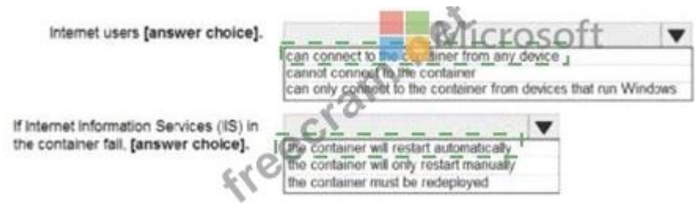

Explanation:

Box 1: can connect to the container from any device

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

Box 2: the container will restart automatically

Docker provides restart policies to control whether your containers start automatically when they exit, or when Docker restarts. Restart policies ensure that linked containers are started in the correct order.

Docker recommends that you use restart policies, and avoid using process managers to start containers.

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

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

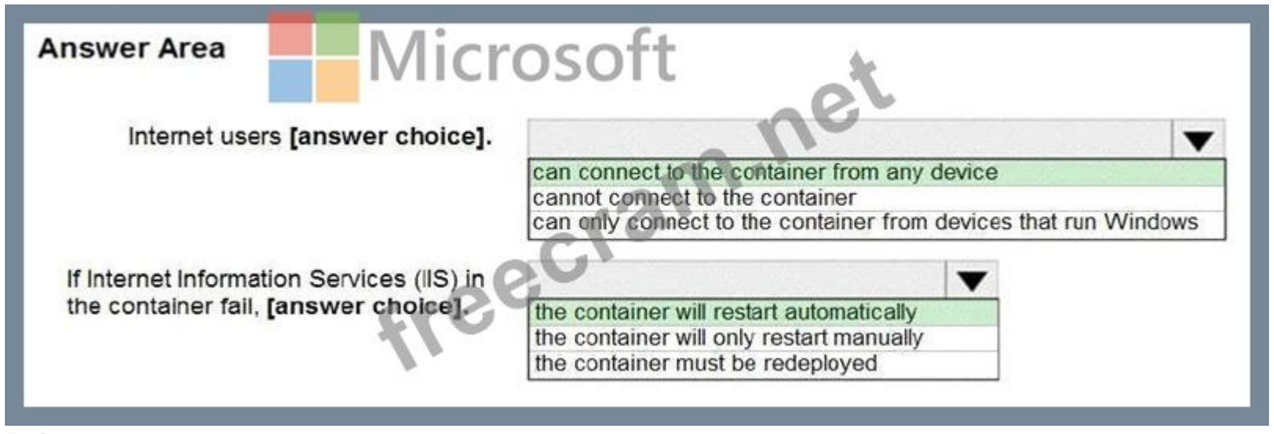

Reference:

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

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

NEW QUESTION: 37

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 VNETl 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: 38

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.

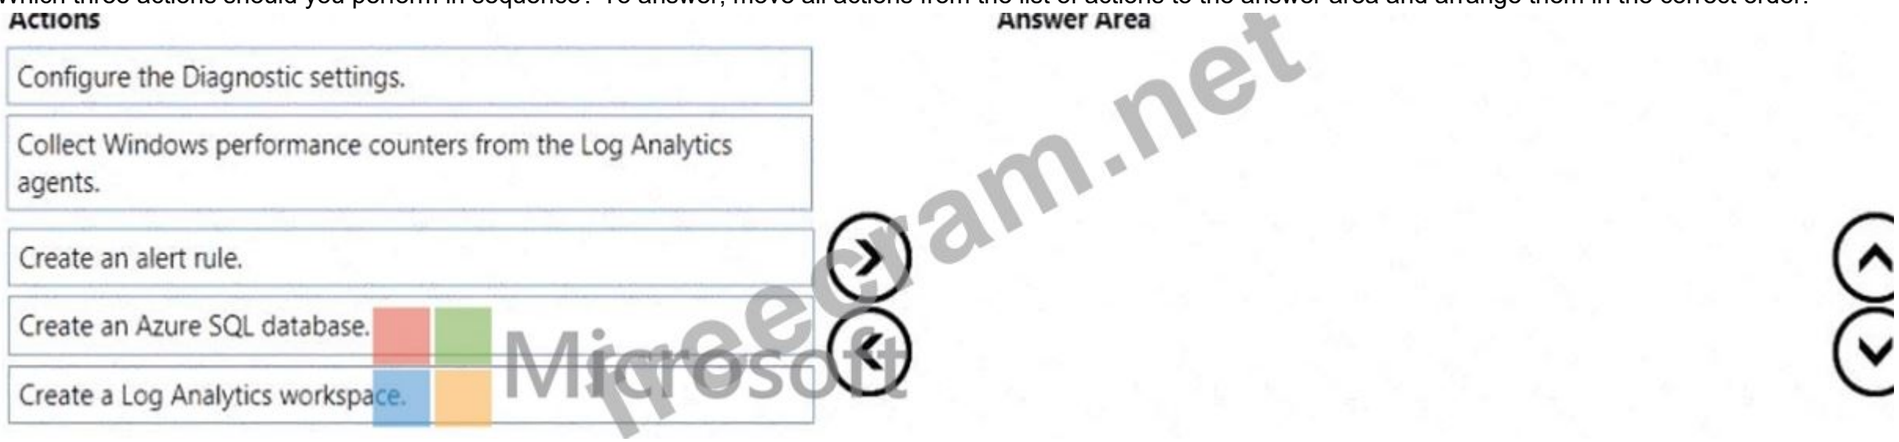

Answer:

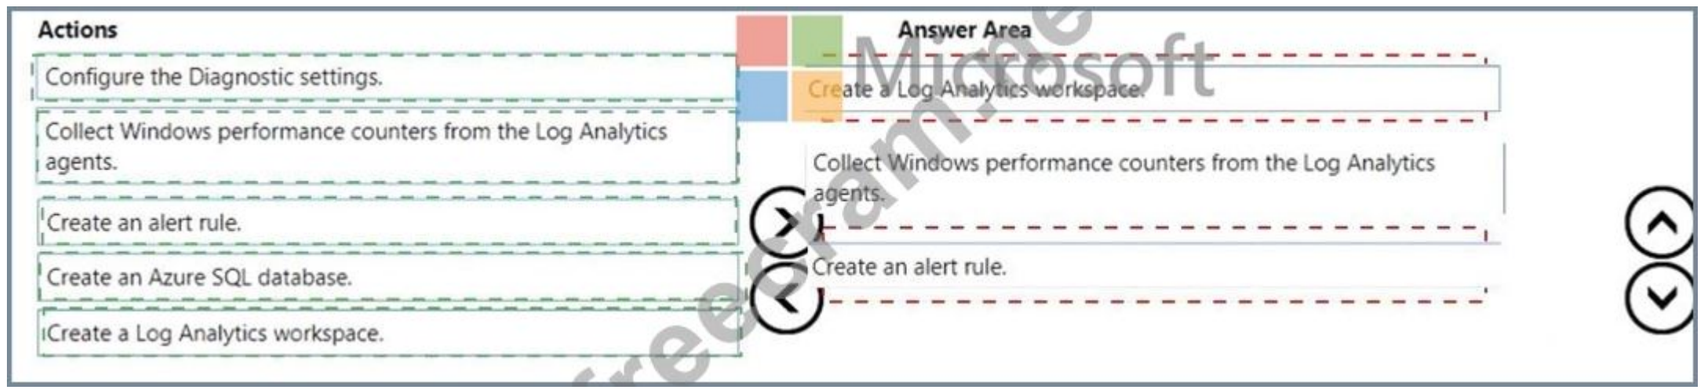

Explanation:

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: 39

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: 40
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: 41
You need to implement Role1.

Which command should you run before you create Role1? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

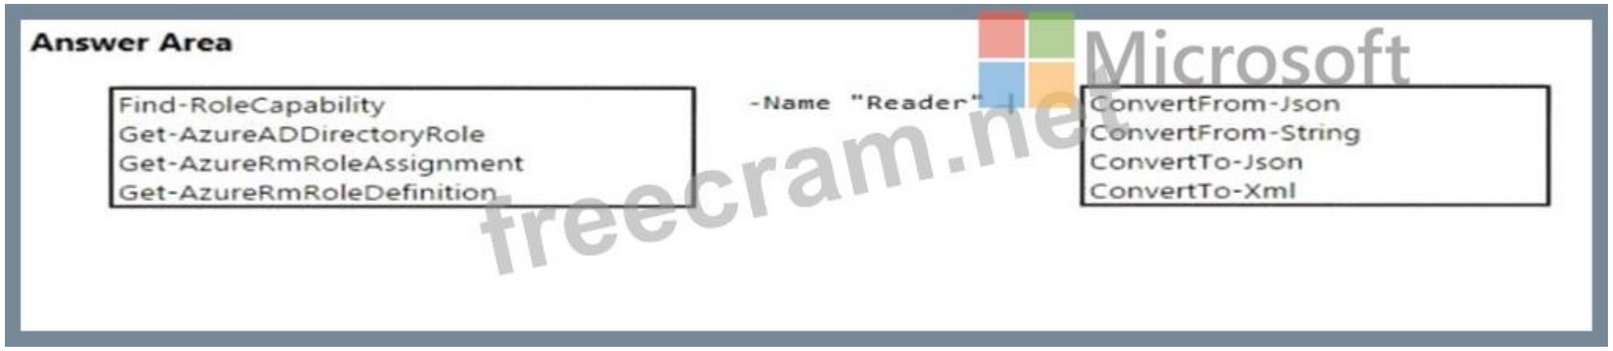

Answer:

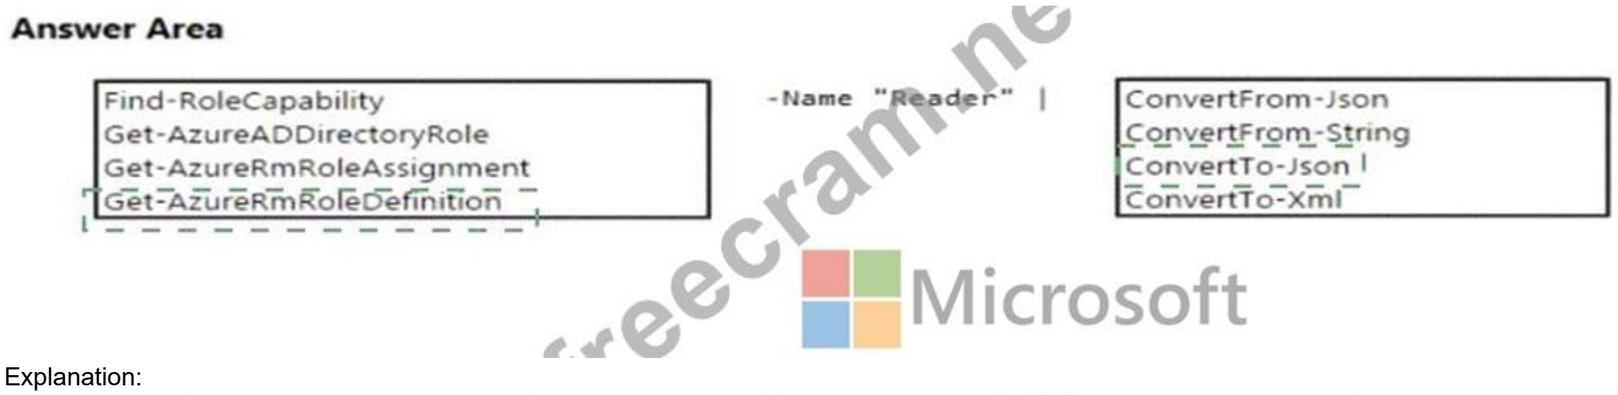

Explanation:

https://docs.microsoft.com/en-us/azure/role-based-access-control/tutorial-custom-role-powershell Get-AzRoleDefinition -Name " Reader " | ConvertTo-Json
https://docs.microsoft.com/en-us/powershell/module/az.resources/get-azroledefinition?view=azps-5.9.0
https://docs.microsoft.com/en-us/azure/role-based-access-control/tutorial-custom-role-powershell
https://docs.microsoft.com/en-us/powershell/module/microsoft.powershell.utility/convertto-json?view=powershell-7.1
https://docs.microsoft.com/en-us/powershell/module/azuread/get-azureaddirectoryrole?view=azureadps-2.0

NEW QUESTION: 42
You plan to move a distributed on-premises app named App1 to an Azure subscription.
After the planned move, App1 will be hosted on several Azure virtual machines.

You need to ensure that App1 always runs on at least eight virtual machines during planned Azure maintenance.
What should you create?
A. one virtual machine scale set that has 10 virtual machines instances
B. one Availability Set that has three fault domains and one update domain
C. one Availability Set that has 10 update domains and one fault domain
D. one virtual machine scale set that has 12 virtual machines instances
Answer: (SHOW ANSWER)
A virtual machine scale set is a group of identical virtual machines that are centrally managed, configured, and updated1. A virtual machine scale set can automatically increase or decrease the number of virtual machine instances in response to demand or a defined schedule2. A virtual machine scale set also provides high availability and fault tolerance by distributing the virtual machine instances across multiple fault domains and update domains3.
A fault domain is a logical group of underlying hardware that share a common power source and network switch. A fault domain can fail due to hardware or software failures, power outages, or network interruptions4. A virtual machine scale set can have up to five fault domains in a region.
An update domain is a logical group of underlying hardware that can undergo maintenance or be rebooted at the same time. An update domain can be affected by planned events, such as OS updates, application updates, or configuration changes4. A virtual machine scale set can have up to 20 update domains in a region.
By creating a virtual machine scale set that has 10 virtual machine instances, you can ensure that App1 always runs on at least eight virtual machines during planned Azure maintenance. This is because the default configuration of a virtual machine scale set is to have five fault domains and five update domains. This means that at any given time, only one fault domain or one update domain can be unavailable due to maintenance or failure. Therefore, at least eight out of 10 virtual machine instances will be available to run App1.
An availability set is another option for providing high availability and fault tolerance for your virtual machines. An availability set is a logical grouping of two or more virtual machines that are deployed across multiple fault domains and update domains. However, an availability set does not provide automatic scaling of resources or load balancing of traffic. You need to manually create and manage the number of virtual machine instances in an availability set.
Therefore, a virtual machine scale set is a better option than an availability set for your scenario. To create a virtual machine scale set, you can follow these steps:
Sign in to the Azure portal.
Select Create a resource > Compute > Virtual machine scale set.
On the Basics tab, enter a name for your scale set, select your subscription and resource group, select Windows Server 2019 as the image type, and enter a username and password for the administrator account.
On the Instance details tab, select the region where you want to deploy your scale set, select the size of the virtual machine instances, and enter 10 as the initial instance count.
On the Scaling tab, configure the scaling policy for your scale set based on metrics or schedule.
On the Load balancing tab, configure the load balancer for your scale set to distribute traffic across the instances.
On the Management tab, configure the diagnostics settings, automatic OS upgrades, extensions, and backup options for your scale set.
On the Advanced tab, configure the availability zone, proximity placement group, accelerated networking, host group, and custom script extension options for your scale set.
On the Tags tab, optionally add tags to your scale set resources.
On the Review + create tab, review your settings and select Create.

NEW QUESTION: 43
You have an Azure Storage account named storage1 that uses Azure Blob storage and Azure File storage.
You need to use AzCopy to copy data to the blob storage and file storage in storage1.
Which authentication method should you use for each type of storage? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

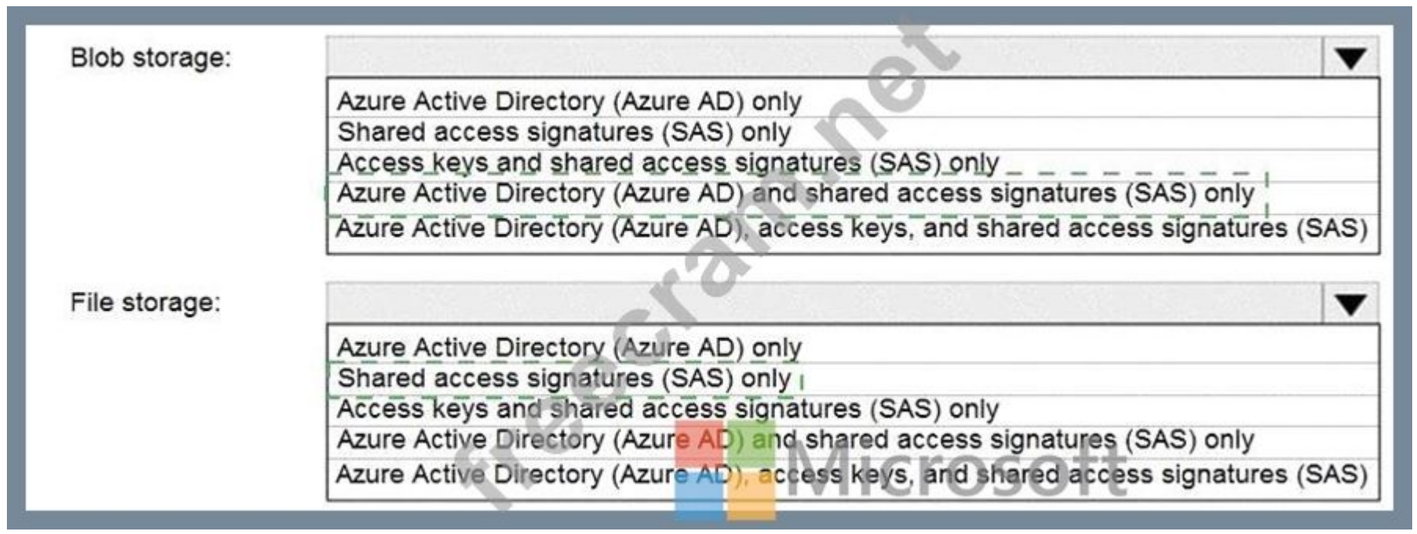

Explanation:

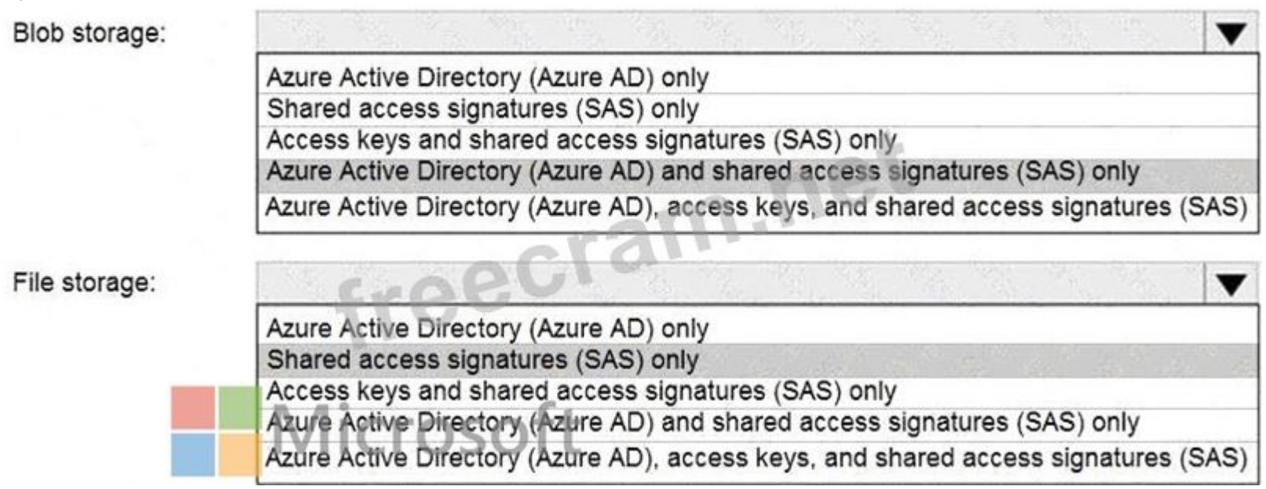

You can provide authorization credentials by using Azure Active Directory (AD), or by using a Shared Access Signature (SAS) token.
Box 1:

Both Azure Active Directory (AD) and Shared Access Signature (SAS) token are supported for Blob storage.
Box 2:
Only Shared Access Signature (SAS) token is supported for File storage.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10

NEW QUESTION: 44
You have an Azure subscription that contains an Azure SQL database named DB1.
You plan to use Azure Monitor to monitor the performance of DB1. You must be able to run queries to analyze log data.
Which destination should you configure in the Diagnostic settings of DB 1?
A. Send to a Log Analytics workspace.
B. Archive to a storage account.
C. Stream to an Azure event hub.
Answer: (SHOW ANSWER)
According to the Microsoft documentation, Azure Monitor collects and analyzes monitoring data from Azure resources, including Azure SQL databases. You can use Azure Monitor to monitor the performance of DB1 and run queries to analyze log data.
To use Azure Monitor, you need to configure the diagnostic settings of DB1, which define the sources and destinations of the monitoring data. The sources are the types of metric and log data to send to the destinations, such as SQLInsights, Errors, Blocks, Deadlocks, etc. The destinations are one or more locations where you want to send the monitoring data, such as a Log Analytics workspace, a storage account, or an event hub.
A Log Analytics workspace is a unique environment for Azure Monitor log data. Each workspace has its own data repository and configuration, and data sources and solutions are configured to store their data in a particular workspace. You can use a Log Analytics workspace to run queries on the log data collected from DB1 and other resources using the Kusto query language. You can also create alerts, dashboards, and workbooks based on the log data in the workspace.
A storage account is a place where you can store large amounts of unstructured data, such as files, blobs, queues, tables, and disks. You can use a storage account to archive the monitoring data from DB1 for long-term retention or backup purposes. However, you cannot run queries on the log data in a storage account directly. You would need to use another tool or service to analyze the log data in a storage account.
An event hub is a service that enables you to ingest and process large volumes of streaming data from multiple sources. You can use an event hub to stream the monitoring data from DB1 to other applications or services that can consume and analyze the data in real time. However, you cannot run queries on the log data in an event hub directly. You would need to use another tool or service to analyze the log data in an event hub.

NEW QUESTION: 45
You have an Azure subscription named Subscription1 that contains the storage accounts shown in the following table:

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: 46
You have an on-premises server that contains a folder named D:\Folder1.
You need to copy the contents of D:\Folder1 to the public container in an Azure Storage account named contoso data.
Which command should you run?
A. https://contosodata.blob.core.windows.net/public
B. azcopy sync D:\folder1 https://contosodata.blob.core.windows.net/public--snapshot
C. azcopy copy D:\folder1 https://contosodata.blob.core.windows.net/public--recursive
D. az storage blob copy start-batch D:\Folder1 https://contosodata.blob.core.windows.net/public
Answer: (SHOW ANSWER)
The azcopy copy command copies a directory (and all of the files in that directory) to a blob container. The result is a directory in the container by the same name.
Reference:
https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs
https://docs.microsoft.com/en-us/azure/storage/common/storage-ref-azcopy-copy

NEW QUESTION: 47
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: 48

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.

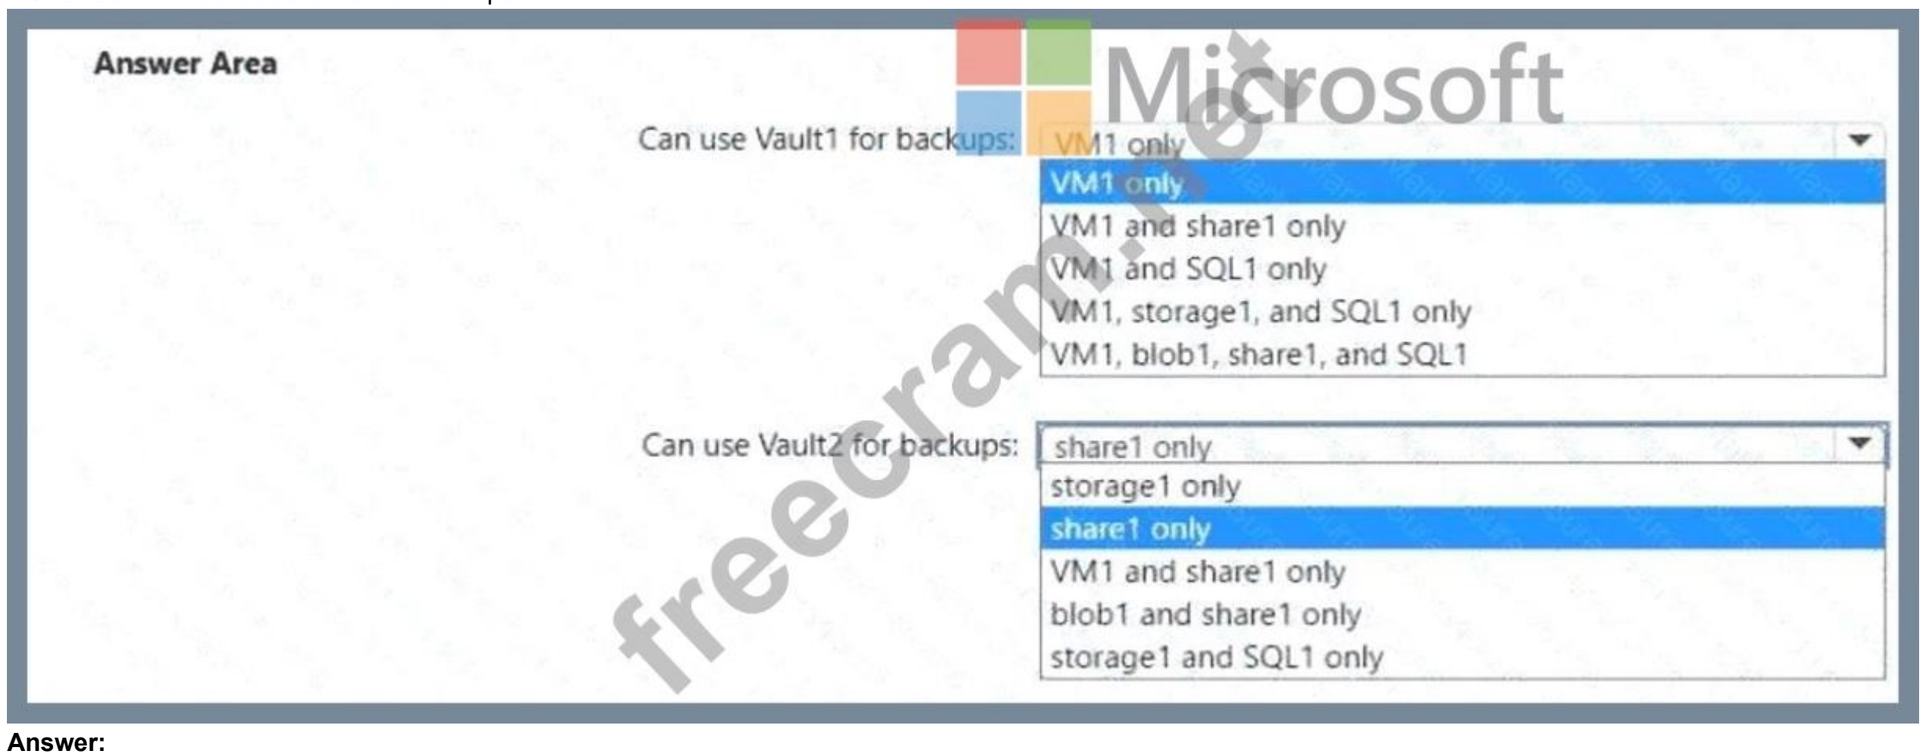

Answer:

Answer Area

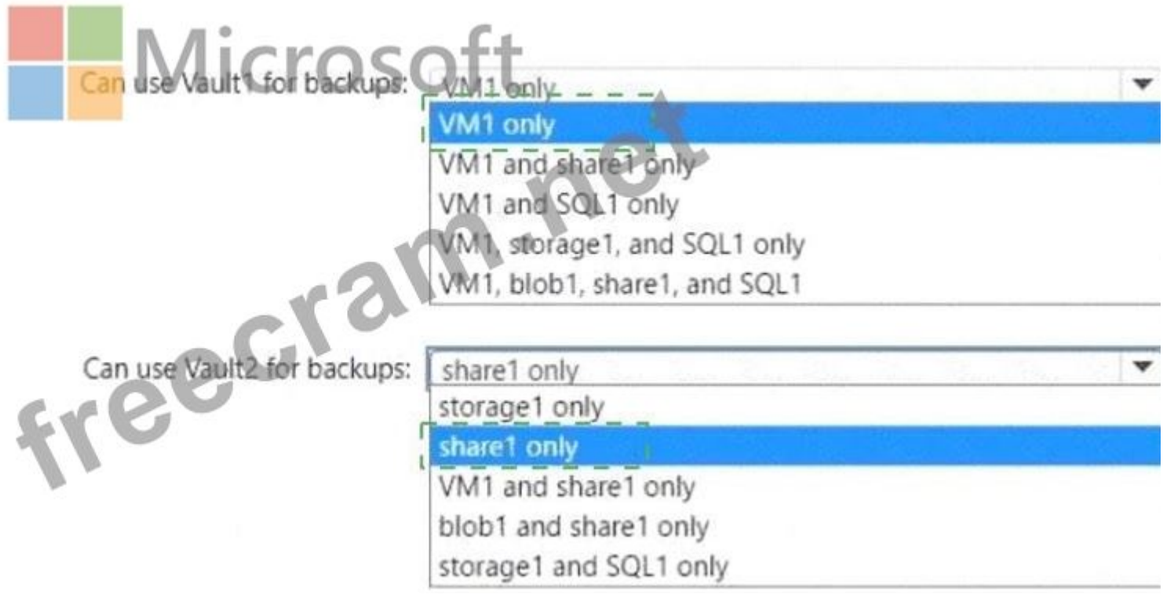

Explanation:

Answer Area

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: 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 an app named App1 that is installed on two Azure virtual machines named VM1 and VM2.
Connections to App1 are managed by using an Azure Load Balancer.
The effective network security configurations for VM2 are shown in the following exhibit.

You discover that connections to App1 from 131.107.100.50 over TCP port 443 fail. You verify that the Load Balancer rules are configured correctly.
You need to ensure that connections to App1 can be established successfully from 131.107.100.50 over TCP port 443.
Solution: You modify the priority of the Allow_131.107.100.50 inbound security rule.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In this scenario, you are troubleshooting inbound network traffic to VM2 where connections from the public IP 131.107.100.50 over TCP port 443 are failing, despite the Azure Load Balancer configuration being correct. The issue lies in the Network Security Group (NSG) rules applied to VM2-NIC1.
Understanding NSG Rule Processing
According to Microsoft Azure Administrator documentation (AZ-104 Study Guide, "Secure network traffic" section):

"Network security group rules are processed in priority order, from the lowest number to the highest number.
As soon as a rule matches the traffic, rule processing stops."
This means that once a matching rule is found (based on source, destination, port, and protocol), Azure will not evaluate further rules.
Analysis of the Exhibit
From the NSG configuration shown:
Priority
Name
Port
Protocol
Source
Destination
Action
100
Allow_131.107.100.50
443
TCP
131.107.100.50
VirtualNetwork
Allow
200
Block_All_Other443
443
TCP
Any
Any
Deny
65000
AllowVNetInBound
Any
Any
VirtualNetwork
VirtualNetwork
Allow
65001
AllowAzureLoadBalancerInBound
Any
Any
AzureLoadBalancer
Any
Allow
65500

DenyAllInBound
Any
Any
Any
Any
Deny
At first glance, Rule 100 should allow traffic from 131.107.100.50 on port 443, but it specifies Destination:
VirtualNetwork, which applies only if the destination IP is part of the same virtual network (VNet).
However, in this case, the traffic originates from an external public IP address (131.107.100.50) trying to reach VM2 through the Azure Load Balancer, not directly through the virtual network.
Because of this, the destination does not match the "VirtualNetwork" scope, and the rule is effectively skipped. The next applicable rule (priority 200) explicitly denies all traffic on port 443. Therefore, the connection fails.
Proposed Solution: Modify the Priority
The question's solution proposes modifying the priority of the Allow_131.107.100.50 rule.
If you adjust the rule's priority and ensure that it is processed before the "Block_All_Other443" deny rule (priority 200) and correct its destination to Any, then the traffic from 131.107.100.50 will be evaluated first and allowed before the deny rule triggers.
Hence, by modifying the priority or adjusting the target scope, the traffic will successfully match the allow rule.
Official Documentation Reference
From Microsoft Learn - Network Security Groups Overview:
"Rules are processed in priority order. Once a rule matches the traffic, processing stops."
"If you have a rule that denies specific traffic with a higher priority, any allow rules with a lower priority won' t be processed." Conclusion By modifying the priority of the Allow_131.107.100.50 rule (or ensuring it appears before the deny rule in evaluation order), inbound traffic from that specific IP over port 443 will be allowed before the deny rule applies.
Therefore, the proposed solution meets the goal.
Final Verified Answer: A. Yes
Modifying the priority of the "Allow_131.107.100.50" rule ensures it is processed before the "Block_All_Other443" deny rule, allowing the required inbound traffic to reach the VM successfully.

NEW QUESTION: 50

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: 51
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.

Create a Bastion

Basics | Tags | Advanced | Review + create

Bastion allows web based RDP access to your vnet VM. Learn more

Project details

Subscription * MSDN Platforms

Resource group * RG1

Create new

Instance details

Name * Bastion1

Virtual network * VNet1

Create new

Subnet * AzureBastionSubnet (10.0.2.0/24)

Manage subnet configuration

Public IP address

Public IP address * Create new / Use existing

Public IP address name * VNet1-ip

Public IP address SKU Standard

Assignment Dynamic / Static

Review + create | Previous | Next : Tags > | Download a template for automation

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

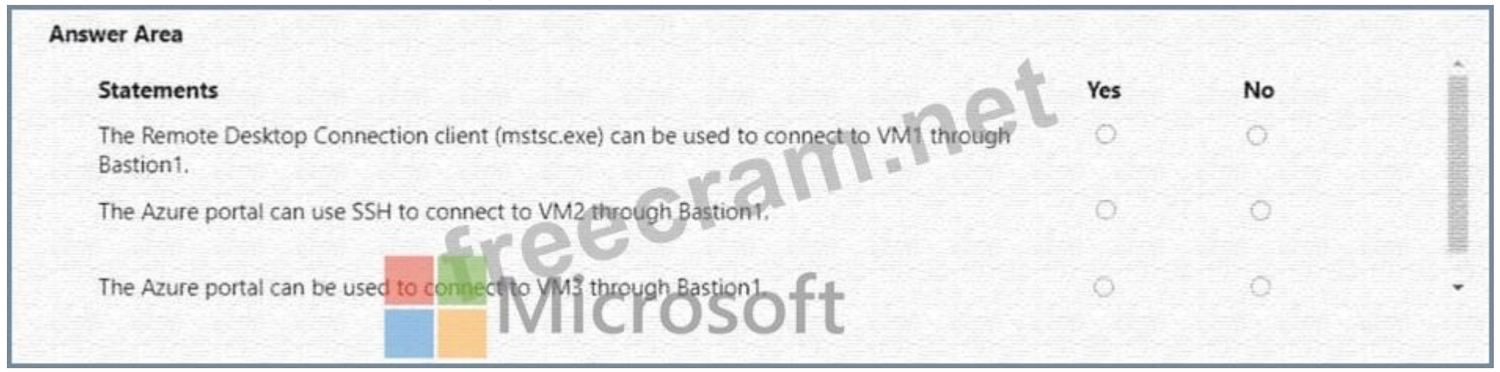

Answer:

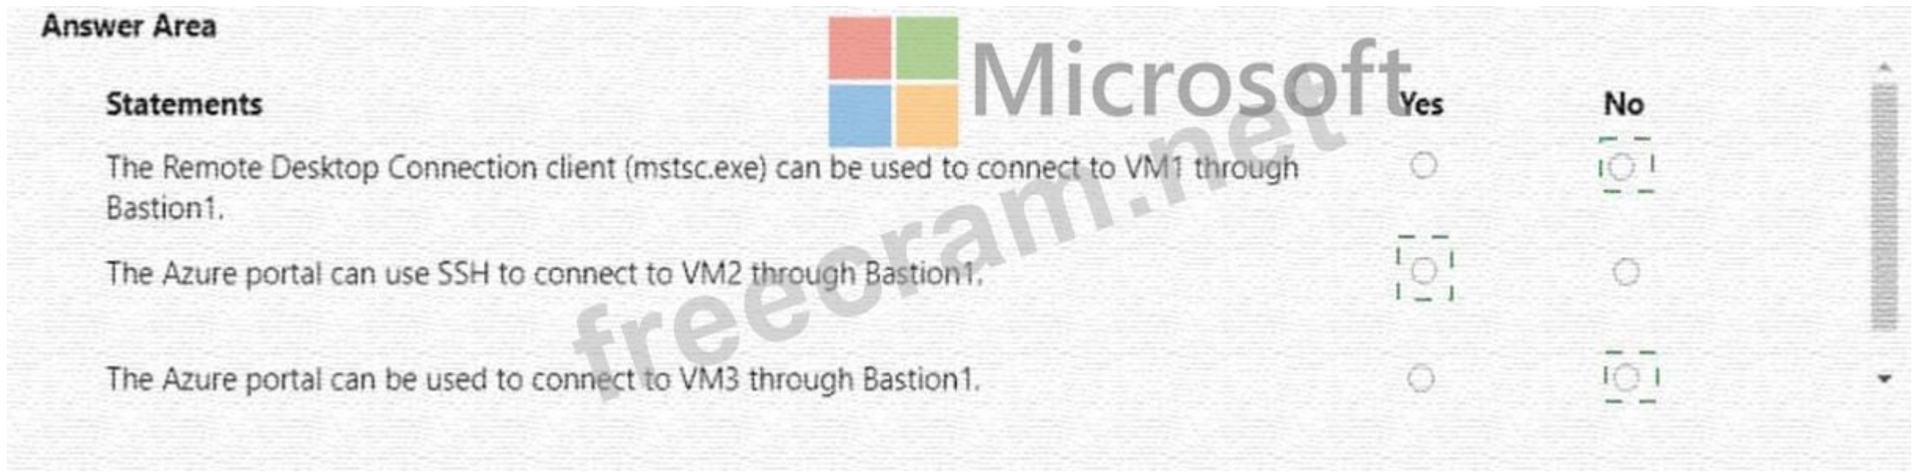

Explanation:
No
Yes
No
Azure Bastion is a fully managed platform-as-a-service (PaaS) solution that allows secure RDP (for Windows VMs) and SSH (for Linux VMs) connections directly through the Azure portal - without exposing public IP addresses. It provides browser-based connectivity over TLS from the 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: 52
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: 53
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: B (LEAVE A REPLY)
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: 54
You have an Azure subscription that contains the resources in the following table.

Name	Type	Azure region	Resource group
VNet1	Virtual network	West US	RG2
VNet2	Virtual network	West US	RG1
VNet3	Virtual network	East US	RG1
NSG1	Network security group (NSG)	East US	RG2

To which subnets can you apply NSG1?
A. the subnets on VNet1 only
B. the subnets on VNet2 only
C. the subnets on VNet3 only
D. the subnets on VNet2 and VNet3 only
E. the subnets on VNet1 VNet2, and VNet3
Answer: (SHOW ANSWER)
In Microsoft Azure, Network Security Groups (NSGs) are used to filter network traffic to and from Azure resources in an Azure Virtual Network (VNet). NSGs can be associated with subnets or network interfaces (NICs) but are scoped by region - meaning that an NSG can only be applied to resources within the same Azure region.
From the scenario:
* NSG1 is located in the East US region.
* VNet1 and VNet2 are located in the West US region.
* VNet3 is located in the East US region.
Therefore, based on Microsoft Azure's network management rules, NSG1 can only be applied to subnets or network interfaces in VNets located in the same region (East US). This means NSG1 can only be assigned to VNet3, since both reside in the East US region.
As stated in Microsoft Learn (Azure Networking documentation):
"A network security group (NSG) must exist in the same region as the virtual network or subnet to which it is associated. You cannot link or associate an NSG across regions." This regional boundary ensures proper security control enforcement and routing consistency within Azure's networking infrastructure.

NEW QUESTION: 55
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: 56

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?

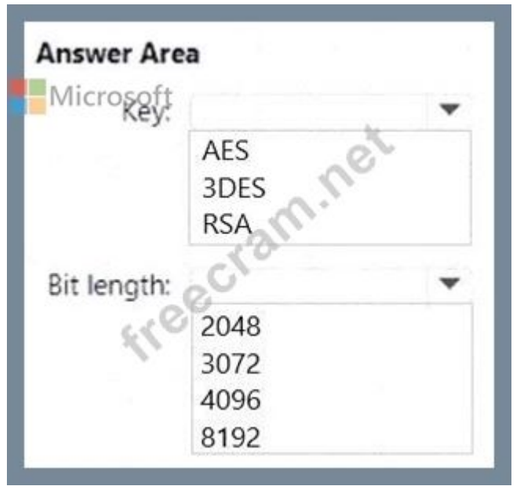

Answer:

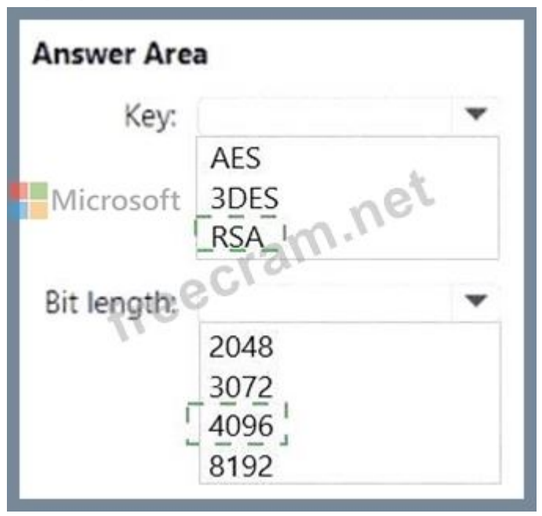

Explanation:

RSA

4096

Key: RSA

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

NEW QUESTION: 57

You have an Azure subscription named Subscroption1.

In Subscription1, you create an alert rule named Alert1.

The Alert1 action group is configured as shown in the following exhibit.

```
PS Azure:\> Get-AzureRmActionGroup

ResourceGroupName: default-activitylogalerts
GroupShortName   : AG1
Enabled          : True
EmailReceivers   : {Action1_-EmailAction-}
SmsReceivers     : {Acrtion_-SMSAction-}
WebhookReceivers : {}
Id               : /subscriptions/a4fde29b-d56a-4f6c-8298-6c53cd0b720c/
resourceGroups/default-activitylogalerts/providers/microsoft.insights/actionGrooups/ActionGroup1
Name             : ActionGroup1
Type             : Microsoft.Insights/ActionGroups
Location         : Global
Tags             : {}
```

Alert1 alert criteria is triggered every minute.

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.

The number of email messages that Alert1 will send in an hour is **[answer choice]**.

- 0
- 4
- 6
- 12
- 60

The number of SMS messages that Alert1 will send in an hour is **[answer choice]**.

- 0
- 4
- 6
- 12
- 60

Answer:

The number of email messages that Alert1 will send in an hour is **[answer choice]**.

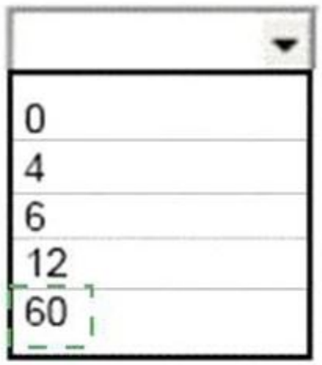

The number of SMS messages that Alert1 will send in an hour is **[answer choice]**.

Explanation:

The number of email messages that Alert1 will send in an hour is **[answer choice]**.

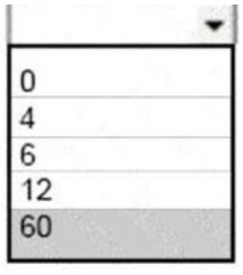

The number of SMS messages that Alert1 will send in an hour is **[answer choice]**.

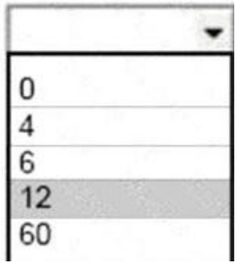

Box 1: 60
One alert per minute will trigger one email per minute.
Box 2: 12
No more than 1 SMS every 5 minutes can be send, which equals 12 per hour.
Note: Rate limiting is a suspension of notifications that occurs when too many are sent to a particular phone number, email address or device. Rate limiting ensures that alerts are manageable and actionable.
The rate limit thresholds are:
SMS: No more than 1 SMS every 5 minutes.
- Voice: No more than 1 Voice call every 5 minutes.
- Email: No more than 100 emails in an hour.
- Other actions are not rate limited.

References:
https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/monitoring-and-diagnostics/monitoring- overview-alerts.md

NEW QUESTION: 58
You have an Azure subscription that contains a virtual machine name VM1.
VM1 has an operating system disk named Disk1 and a data disk named Disk2.
You need to back up Disk2 by using Azure Backup.
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.

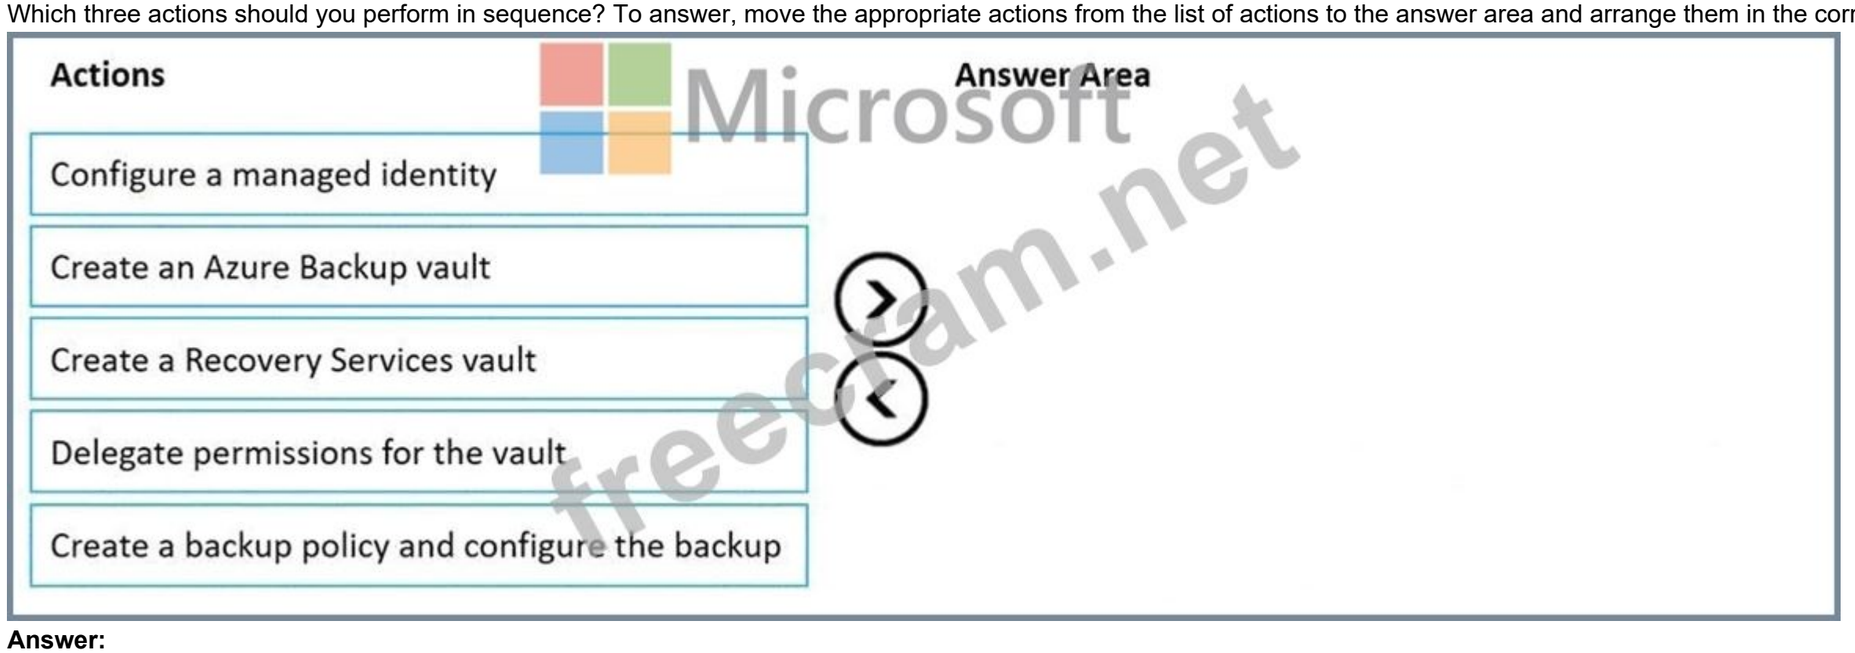

Answer:

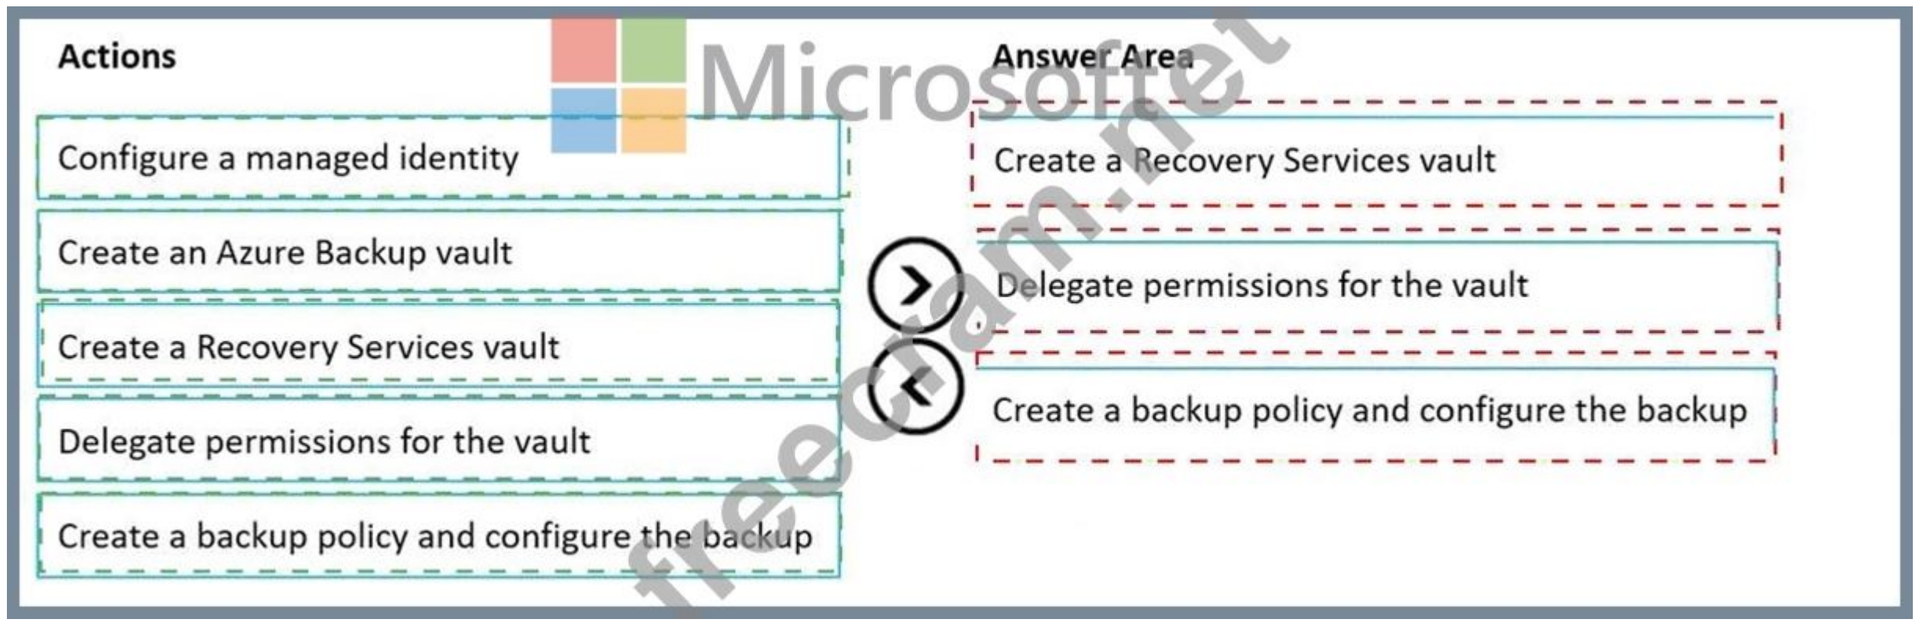

Explanation:

1. Create a Recovery Services vault
2. Delegate permissions for the vault
3. Create a backup policy and configure the backup

When backing up a specific disk (like Disk2) of an Azure virtual machine using Azure Backup, the process involves the use of a Recovery Services vault - not an "Azure Backup vault." The Recovery Services vault is the central management entity for backup policies, recovery points, and monitoring of backup operations.

According to Microsoft Azure Administrator (AZ-104) study documentation and Azure Backup official documentation, the correct sequential procedure is:

1## Create a Recovery Services vault

* A Recovery Services vault stores backup data and recovery points. It is a required component before you can configure backups for any VM disks.
* Extract from documentation:

"A Recovery Services vault is a storage entity in Azure that houses backup data for various workloads such as Azure VMs, SQL databases, and files."

2## Delegate permissions for the vault

* The Azure Backup service needs permission to access and manage the VM's resources during backup operations. By default, this is handled through the Azure Backup service principal.
* For custom scenarios (such as disk-level backup or managed identity use), proper permissions must be granted to allow access to both the vault and the VM's disks.
* Extract from documentation:

"Ensure that the Backup service has the necessary permissions on the vault and virtual machine resource group to configure backups and access data disks."

3## Create a backup policy and configure the backup

* Finally, create a backup policy (daily, weekly, etc.) and apply it to VM1, selecting Disk2 specifically for protection.
* Extract from documentation:

"After the Recovery Services vault is created and permissions are delegated, configure a backup policy that defines the schedule and retention, then enable backup for the selected disks." This process ensures compliance with Azure Backup's architecture and automation framework while using Azure Resource Manager (ARM) authorization flow.

Final Verified Sequence Answer:

* Create a Recovery Services vault
* Delegate permissions for the vault
* Create a backup policy and configure the backup

NEW QUESTION: 59

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

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

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

D. comp2 contoso.com only

E. comp2.contoso.com and comp4.contoso.com only

Answer: (SHOW ANSWER)

NEW QUESTION: 60

You have an Azure virtual machine named VM1 that runs Windows Server. The VM was deployed using default drive settings. You sign in to VM1 as a user named User1 and perform the following actions:

* Create files on drive C.
* Create files on drive D.
* Modify the screen saver timeout.
* Change the desktop background.

You plan to redeploy VM1. Which changes will be lost after you redeploy VM1?

A. the new files on drive D

B. the modified screen saver timeout

C. the new desktop background

D. the new files on drive C

Answer: (SHOW ANSWER)

Topic 4, 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: 61

You have an Azure subscription that contains a user named User1.

You need to ensure that User1 can deploy virtual machines and manage virtual networks. The solution must use the principle of least privilege.

Which role-based access control (RBAC) role should you assign to User1?

A. Contributor

B. Virtual Machine Contributor

C. Owner

D. Virtual Machine Administrator Login

Answer: (SHOW ANSWER)

NEW QUESTION: 62
You have an Azure virtual machine named VM1. VM1 was deployed by using a custom Azure Resource Manager template named ARMljson.
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: B (LEAVE A REPLY)
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: 63
You have a Microsoft Entra tenant that contains the identities shown in the following table.

Name	Type
User1	User
Group1	Security group
Principal1	Service principal

You have a custom security attribute named Attri. You need to assign Attr1 to all the supported identities. To which identities can you assign Attri?
A. User1 only
B. User1 and Principal1 only
C. User1, Group1. and Principal1
D. User1 and Group1 only
E. Group1 only
Answer: (SHOW ANSWER)

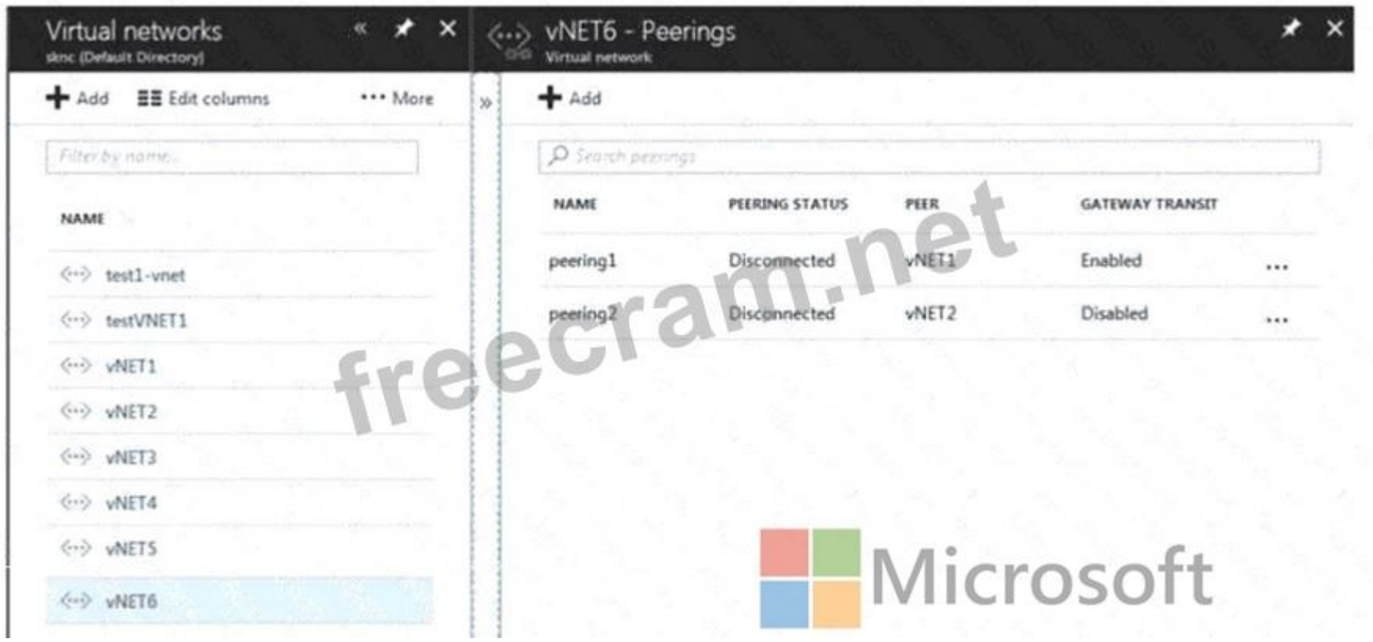

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer:

Answer Area

Hosts on vNET6 can communicate with hosts on **[answer choice]**.

- vNET6 only
- vNET6 and vNET 1 only
- vNET6, vNET1, and vNET2 only
- all the virtual networks in the subscription

To change the status of the peering connection to vNET1 to **Connected**, you must first **[answer choice]**.

- add a service endpoint
- add a subnet
- delete peering1
- modify the address space

Explanation:

In Azure, Virtual Network Peering connects two virtual networks (VNets) and allows traffic between them using private IP addresses, traversing the Microsoft backbone network rather than the public internet. For peering to function correctly, both sides of the connection must have a valid, reciprocal configuration (peering must be created from each VNet and accepted).

From the exhibit:

* vNET6 has two peering configurations:
* peering1 # vNET1 (Gateway Transit = Enabled)
* peering2 # vNET2 (Gateway Transit = Disabled)Both are currently showing a status of

"Disconnected", meaning that the remote peering (from vNET1 and vNET2 back to vNET6) hasn't been set up or accepted yet.

Because the peering links are disconnected, vNET6 cannot communicate with any other VNet-only internal communication within vNET6 is possible.

To fix the peering connection status for vNET1, you cannot directly edit an incomplete or mismatched configuration. Microsoft's official documentation (" Azure Virtual Network peering - requirements and constraints ") states that you must delete and recreate the peering if the configuration is invalid or disconnected. Therefore, to establish a working connection to vNET1, you must first delete peering1 and then reconfigure it correctly on both VNets.

NEW QUESTION: 65

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: 66

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. 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)

D: Seamless SSO works with any method of cloud authentication - Password Hash Synchronization or Pass- through Authentication, and can be enabled via Azure AD Connect.
B: You can gradually roll out Seamless SSO to your users. You start by adding the following Azure AD URL to all or selected users' Intranet zone settings by using Group Policy in Active Directory:
https://autologon.
microsoftazuread-sso.com

NEW QUESTION: 67

You have two Azure App Service web apps named App1 and App2. Each web app has a production deployment slot and a test deployment slot.
The Backup Configuration settings for the web app production slots are shown in the following table.

Web app name	Backup Every	Start backup schedule from	Retention (Days)	Keep at least one backup
App1	1 Days	January 6, 2024	0	Yes
App2	1 Days	January 6, 2024	30	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:

Explanation:

NEW QUESTION: 68
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: 69
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:

Answer Area

Statements	Yes	No
You can create a virtual network in Subscription1.	○	●
You can create a virtual machine in Subscription2.	●	○
You can move Subscription1 to ManagementGroup11.	●	○

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: 70

You have a Microsoft Entra tenant.
You plan to perform a bulk import of users by using the Azure portal.
You need to ensure that imported user objects are added automatically as the members of a specific group based on each user ' s department. The solution must minimize administrative effort.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

NEW QUESTION: 71
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.

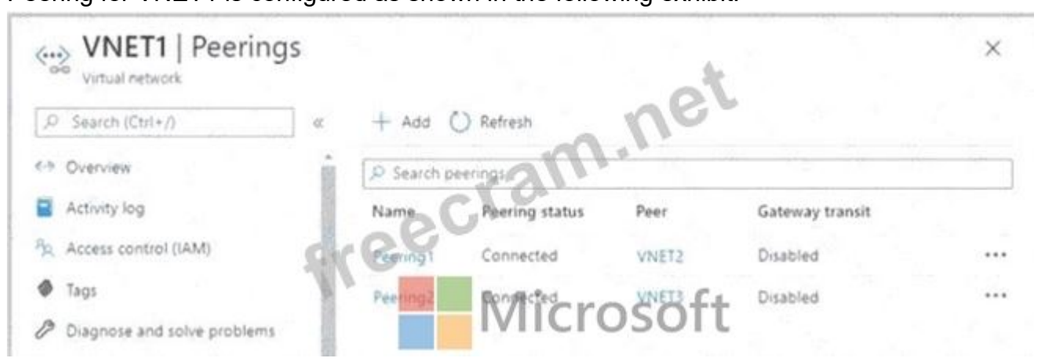

Answer Area

Answer:

Answer Area

Explanation:

Answer Area

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: 72
You have an Azure subscription that contains a storage account named storage1. The storage1 account contains a blob that stores images.
Client access to storage1 is granted by using a shared access signature (SAS).
You need to ensure that users receive a warning message when they generate a SAS that exceeds a seven-day time period.
What should you do for storage1?
A. Add a lifecycle management rule.
B. Set Allow recommended upper limit for shared access signature (SAS) expiry interval to Enabled.
C. Configure an alert rule.
D. Enable a read-only lock.
Answer: (SHOW ANSWER)
Azure Storage provides a built-in security setting that warns users when creating a SAS token with an expiry exceeding seven days.
When Allow recommended upper limit for SAS expiry interval is enabled:
* Users receive a warning message
* SAS creation is not blocked
* Enforces Microsoft's security best practice
Lifecycle rules manage blob data, not access. Alerts cannot intercept SAS creation. Resource locks do not affect SAS behavior.
Microsoft documentation states:
"Enabling the recommended SAS expiry interval warns users when the expiry exceeds seven days."

NEW QUESTION: 73
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.

Actions

- Mount a VHD.
- Copy the files by using File Explorer.
- Download and run a script.
- Select a restore point.
- Copy the files by using AZCopy.
- From the Azure portal, click **Restore VM** from the vault.
- From the Azure portal, click **File Recovery** from the vault.

Answer Area

Answer:

Actions

- Mount a VHD.
- Copy the files by using File Explorer.
- Download and run a script.
- Select a restore point.
- Copy the files by using AZCopy.
- From the Azure portal, click **Restore VM** from the vault.
- From the Azure portal, click **File Recovery** from the vault.

Answer Area

- From the Azure portal, click **File Recovery** from the vault.
- Select a restore point.
- Download and run a script.
- Copy the files by using AZCopy.

Explanation:

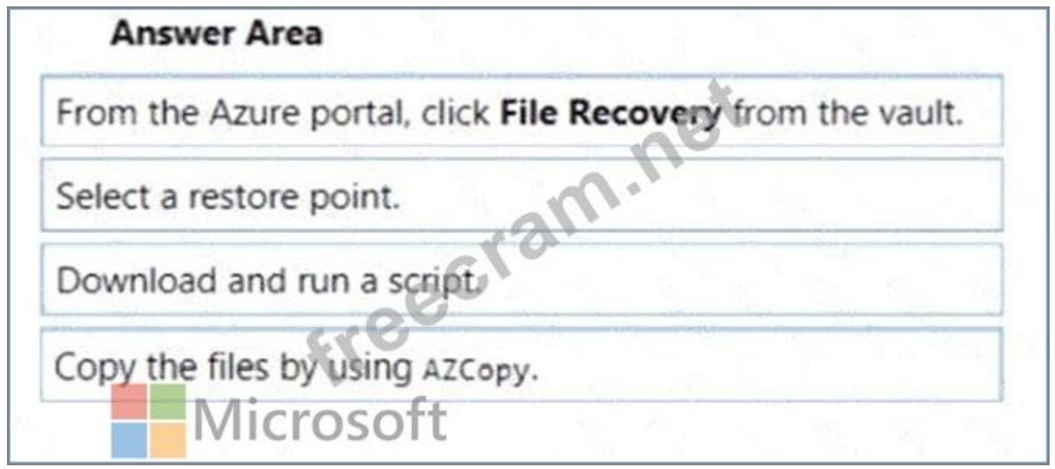

To restore files or folders from the recovery point, go to the virtual machine and choose the desired recovery point.
Step 0. In the virtual machine ' s menu, click Backup to open the Backup dashboard.
Step 1. In the Backup dashboard menu, click File Recovery.
Step 2. From the Select recovery point drop-down 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

NEW QUESTION: 74
Your on-premises network contains a VPN gateway.
You have an Azure subscription that contains the resources shown in the following table.

Name	Type	Description
vgw1	Virtual network gateway	Gateway for Site-to-Site VPN to the on-premises network
storage1	Storage account	Standard performance tier
Vnet1	Virtual network	Enabled forced tunneling
VM1	Virtual machine	Connected to Vnet1

You need to ensure that all the traffic from VM1 to storage! travels across the Microsoft backbone network.
What should you configure?
A. private endpoints
B. Azure Firewall
C. Azure AD Application Proxy
D. Azure Peering Service
Answer: (SHOW ANSWER)
In this scenario, VM1 is connected to VNet1, which has forced tunneling enabled. Forced tunneling ensures that all outbound traffic from the virtual network is directed to the on-premises network through the VPN gateway (vgw1). However, traffic from Azure resources (like VM1) to Azure Storage (storage1) would normally travel over the public internet unless special configuration is applied.
To ensure that the traffic from VM1 to storage1 stays within the Microsoft backbone network (i.e., the private Azure network infrastructure rather than the internet), Microsoft recommends using Azure Private

Endpoints.
A Private Endpoint is a network interface that connects you privately and securely to a service powered by Azure Private Link. When a private endpoint is created for an Azure Storage account, the service gets assigned a private IP address from your virtual network, making it accessible only within that VNet (and via connected VNets through VPN or ExpressRoute).
This ensures that traffic between VM1 and storage1 never traverses the public internet. Instead, it stays entirely within Microsoft's internal network, thereby providing:
* Enhanced security (no exposure to the public internet).
* Improved performance and reliability due to backbone routing.
* Compliance with forced tunneling policies by keeping Azure service traffic within the private network.
Other options are incorrect:
* Azure Firewall (B): Controls traffic but doesn't guarantee Microsoft backbone routing.
* Azure AD Application Proxy (C): Used for publishing on-premises applications, not for controlling storage network paths.
* Azure Peering Service (D): Provides optimized routing for customer internet connectivity, not intra- Azure traffic routing.
Thus, according to Microsoft Learn and AZ-104 Administrator Study Guide, the correct configuration to ensure all traffic from VM1 to storage1 stays on the Microsoft backbone network is the use of Private Endpoints.

NEW QUESTION: 75
You have an Azure subscription.
You plan to create a role definition to meet the following requirements:
* Users must be able to view the configuration data of a storage account.
* Users must be able to perform all actions on a virtual network.
* The solution must use the principle of least privilege.
What should you include in the role definition for each requirement? To answer, select the appropriate options in the answer area.

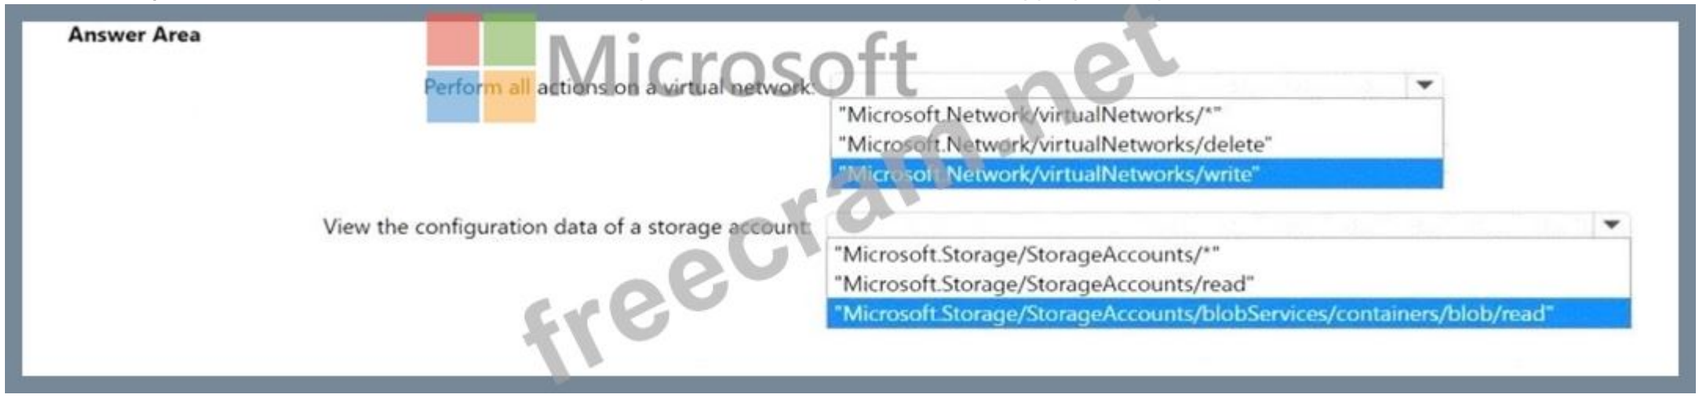

Answer:

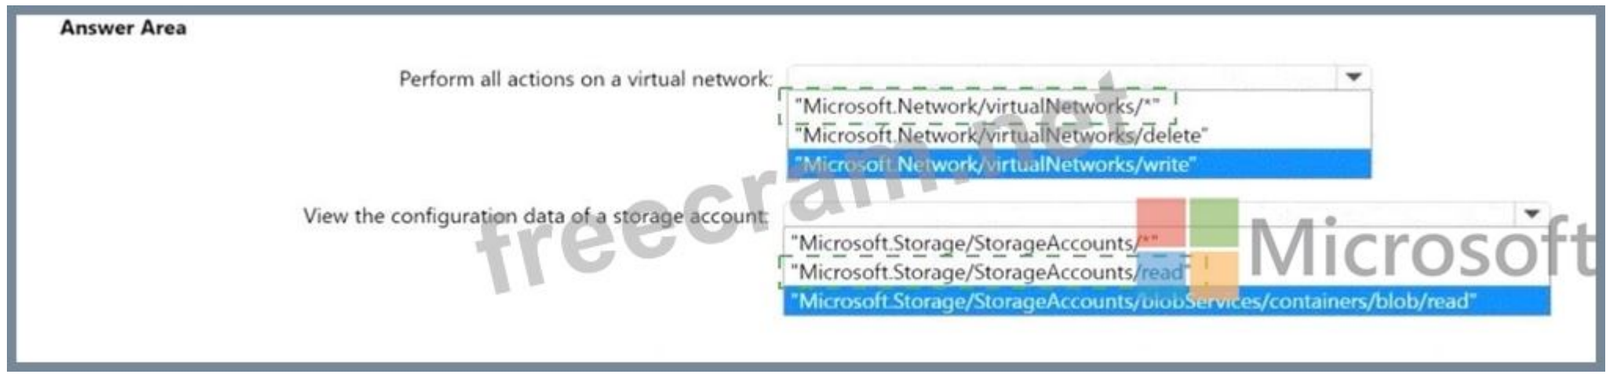

Explanation:
Perform all actions on a virtual network:
"Microsoft.Network/virtualNetworks/*"
View the configuration data of a storage account:
"Microsoft.Storage/StorageAccounts/read"
To perform all actions on a virtual network, you need to use the wildcard () character in the action string, which grants access to all actions that match the string. The action string for virtual networks is "Microsoft. Network/virtualNetworks/". To view the configuration data of a storage account, you need to use the read action substring in the action string, which enables read actions (GET). The action string for storage accounts is "Microsoft.Storage/StorageAccounts/read". References:
https://learn.microsoft.com/en-us/azure/role-based-access-control/role-definitions
https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles

NEW QUESTION: 76
You have five Azure virtual machines that run Windows Server 2016. The virtual machines are configured as web servers.
You have an Azure load balancer named LB1 that provides load balancing services for the virtual machines.
You need to ensure that visitors are serviced by the same web server for each request.
What should you configure?
A. Floating IP (direct server return) to Enabled
B. Idle Time-out (minutes) to 20
C. Protocol to UDP
D. Session persistence to Client IP and Protocol
Answer: (SHOW ANSWER)
https://learn.microsoft.com/en-us/azure/load-balancer/distribution-mode-concepts Session persistence: Client IP and protocol - Traffic from the same client IP and protocol is routed to the same backend instance Azure Load Balancer distributes incoming network traffic among healthy backend pool instances. By default, the load balancer uses a hash-based distribution algorithm based on parameters such as source IP, destination IP, and port. This means that different requests from the same client might be directed to different backend servers.
To ensure that a client is consistently served by the same backend VM (web server) for the duration of its session, Session Persistence (also known as client affinity) must be configured.
As per the Azure Load Balancer official documentation:
"Session persistence specifies how traffic from a client should be handled by the load balancer. When you set session persistence to Client IP or Client IP and Protocol, requests from the same client IP address (and optionally protocol) are directed to the same backend instance." In web server scenarios where session state is maintained in-memory (not in a distributed cache), enabling Client IP and Protocol ensures that users remain connected to the same backend node, preventing issues like session drops or inconsistent experience.
Other options do not meet the requirement:
Floating IP (direct server return): Used for SQL Always On or NVA scenarios, not session affinity.

Idle timeout: Defines TCP connection time-out, not persistence.
UDP protocol: Not relevant for HTTP-based traffic.
Hence, the correct and Microsoft-documented answer is Session persistence = Client IP and Protocol.

NEW QUESTION: 77
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: 78

You have a Microsoft Entra tenant named contoso.onmicrosoft.com that contains the users shown in the following table.

Name	Member of	Role assigned
User1	Group1	None
User2	Group2	None
User3	Group1, Group2	User Administrator

You enable password reset for contoso.onmicrosoftxom as shown in the Password Reset exhibit. (Click the Password Reset tab.)

You configure the authentication methods for password reset as shown in the Authentication Methods exhibit.
(Click the Authentication Methods tab.)

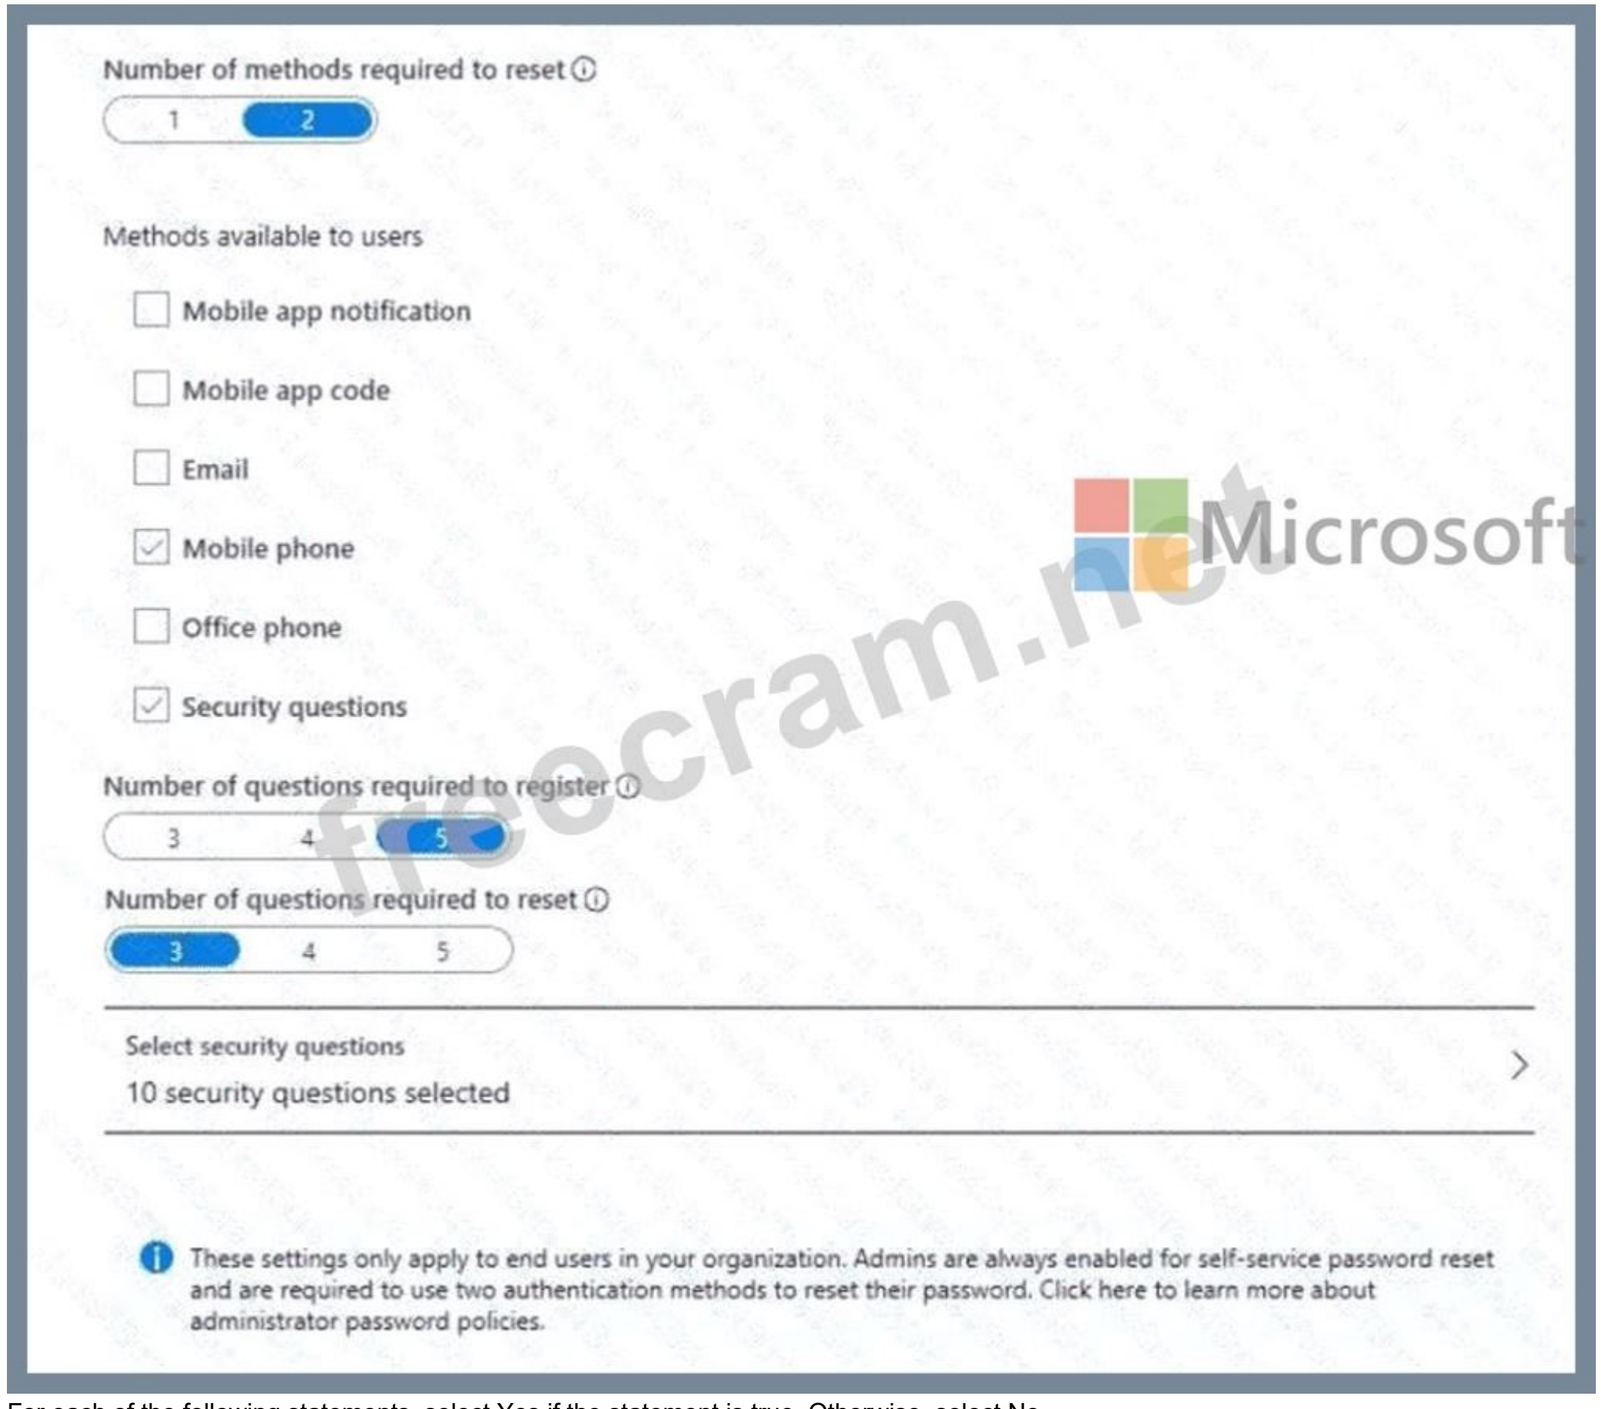

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
After User2 answers three security questions correctly, he can reset his password immediately.	○	○
If User1 forgets her password, she can reset the password by using the mobile phone app.	○	○
User3 can add security questions to the password reset process.	○	○

Answer:

Answer Area

Statements	Yes	No
After User2 answers three security questions correctly, he can reset his password immediately.	○	○
If User1 forgets her password, she can reset the password by using the mobile phone app.	○	○
User3 can add security questions to the password reset process.	○	○

Explanation:

Answer Area

Statements	Yes	No
After User2 answers three security questions correctly, he can reset his password immediately.	○	●
If User1 forgets her password, she can reset the password by using the mobile phone app.	○	●
User3 can add security questions to the password reset process.	●	○

NEW QUESTION: 79

You have an Azure virtual machine named VM1 that runs Windows Server 2019.

You save VM1 as a template named Template1 to the Azure Resource Manager library.

You plan to deploy a virtual machine named VM2 from Template1.

What can you configure during the deployment of VM2?

A. virtual machine size

B. operating system
C. administrator username
D. resource group
Answer: (SHOW ANSWER)
Resource Group is the correct answer: Admin user, password, vm size and os are the part of ARM templates.
But resource group is not hence needs to be mentioned while deployment! Refer below sample ARM template for reference in which all above attributes passed in parameter. https://github.com/Azure/azure-quickstart- templates/blob/master/101-vm-simple-windows/azuredeploy.json

NEW QUESTION: 80
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.

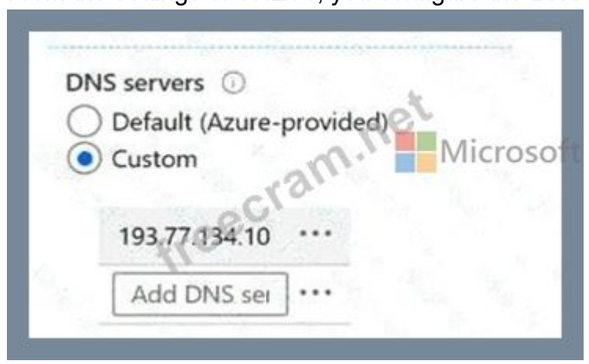

The virtual machines can successfully connect to the DNS server that has an IP address of 192.168.10.15 and the DNS server that has an IP address of 193.77.134.10.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.

	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:

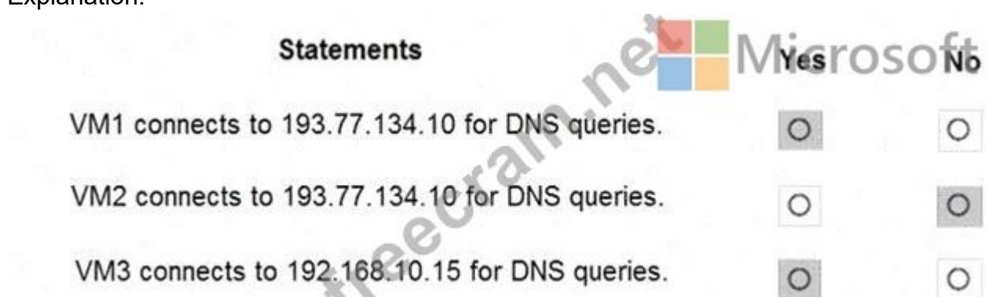

In Azure, DNS resolution order follows a specific hierarchy based on custom DNS configurations at the virtual network (VNet) and network interface (NIC) levels.

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: 81

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

Name	Type	Details
VNet1	Virtual network	*Not applicable*
Subnet1	Subnet	Hosted on VNet1
VM1	Virtual machine	On Subnet1
VM2	Virtual machine	On Subnet1

VM1 and VM2 are deployed from the same template and host line-of-business applications accessed by using Remote Desktop. You configure the network security group (NSG) shown in the exhibit. (Click the Exhibit button.)

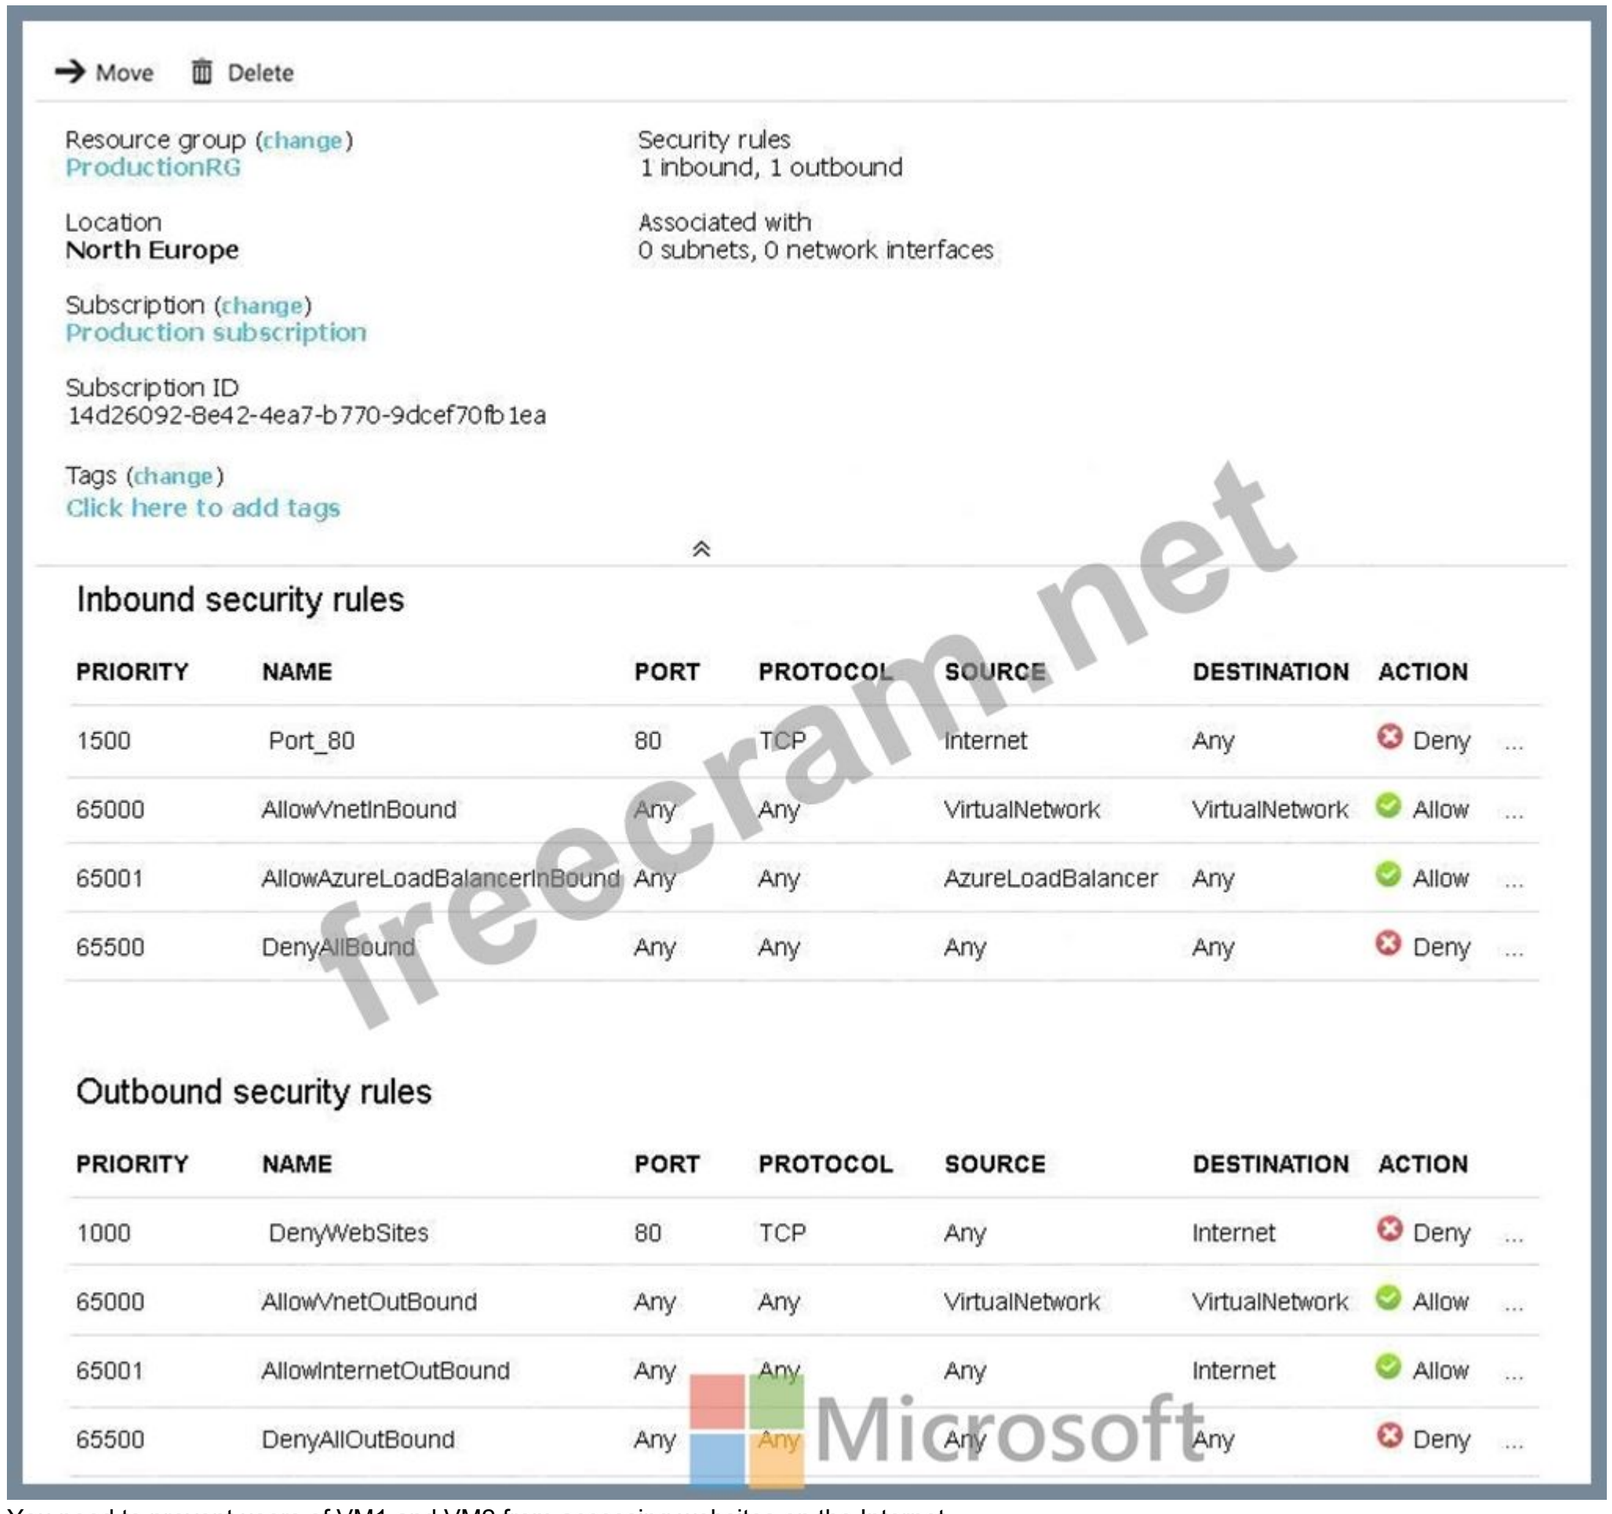

You need to prevent users of VM1 and VM2 from accessing websites on the Internet.

What should you do?

A. Associate the NSG to Subnet1.

B. Disassociate the NSG from a network interface.

C. Change the DenyWebSites outbound security rule.

D. Change the Port_80 inbound security rule

Answer: (SHOW ANSWER)

Outbound rule "DenyWebSites" is setup correctly to block outbound internet traffic over port 80. In the screenshot it states, " Associated with: 0 subnets, 0 NIC ' s " , so you need to associate the NSG to Subnet1.

You can associate or dissociate a network security group from a NIC or Subnet. Reference: https://docs.

microsoft.com/en-us/azure/virtual-network/manage-network-security-group

NEW QUESTION: 82

You plan to create an Azure virtual machine named VM1 that will be configured as shown in the following exhibit.

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

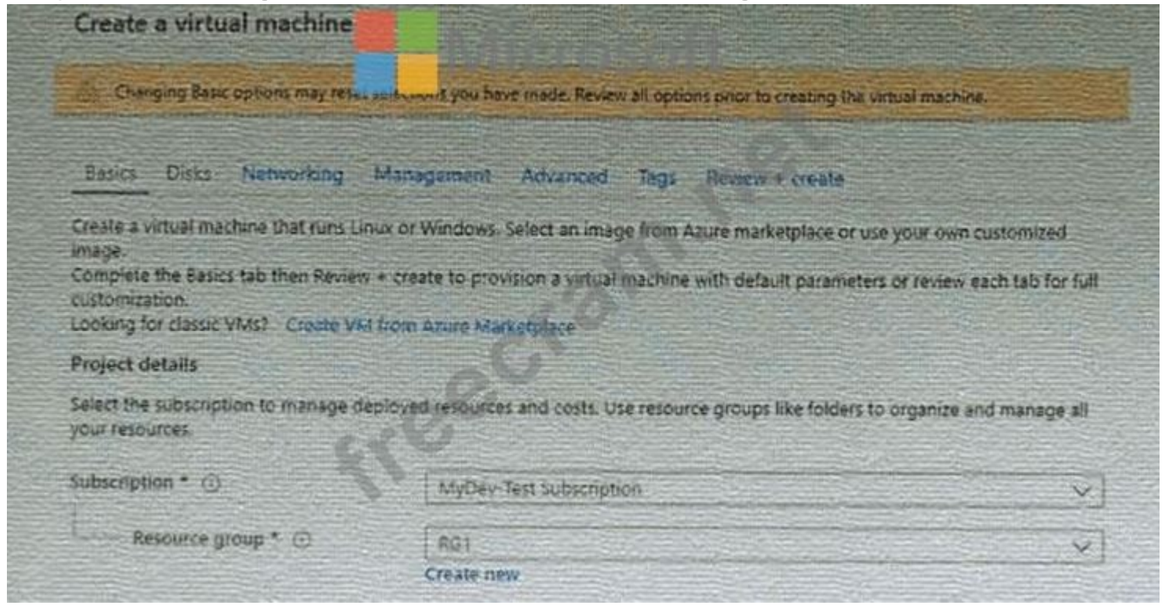

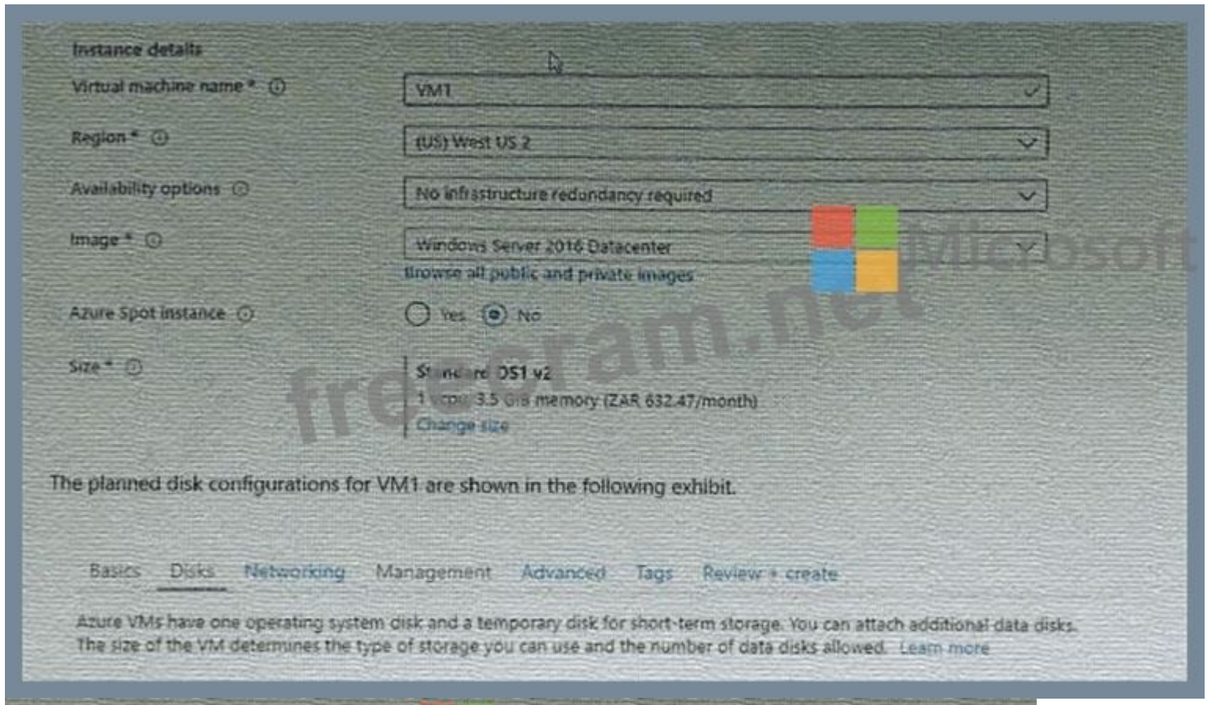

The planned disk configurations for VM1 are shown in the following exhibit.

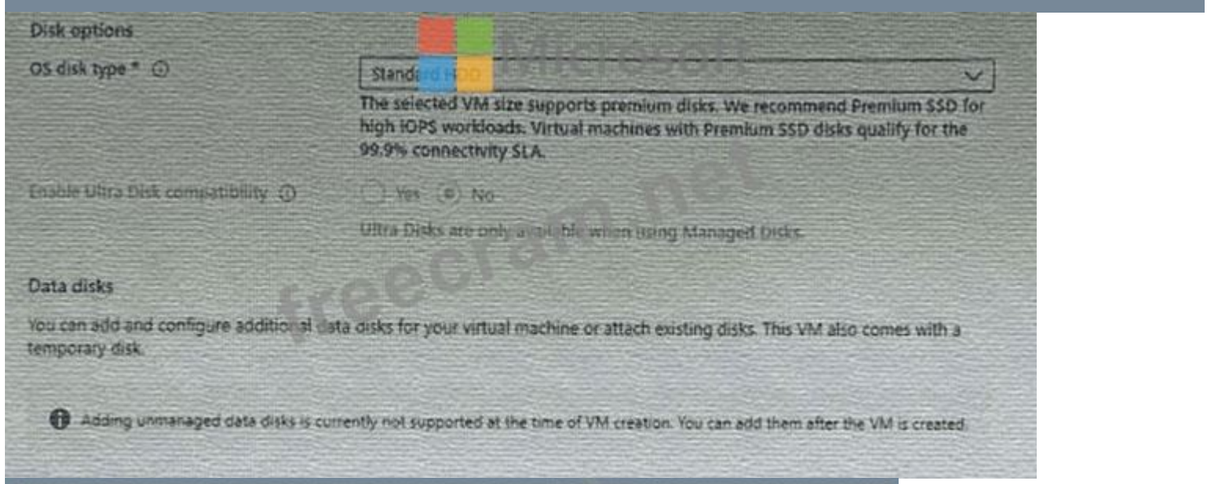

You need to ensure that VM1 can be created in an Availability Zone.

Which two settings should you modify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Use managed disks

B. Availability options
C. OS disk type
D. Size
E. Image
Answer: (SHOW ANSWER)
https://docs.microsoft.com/en-us/azure/site-recovery/move-azure-vms-avset-azone https://docs.microsoft.com /en-us/azure/virtual-machines/windows/create-portal-availability-zone https://docs.microsoft.com/en-us/azure /virtual-machines/manage-availability https://docs.microsoft.com/en-us/azure/availability-zones/az- overview#availability-zones

NEW QUESTION: 83
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.

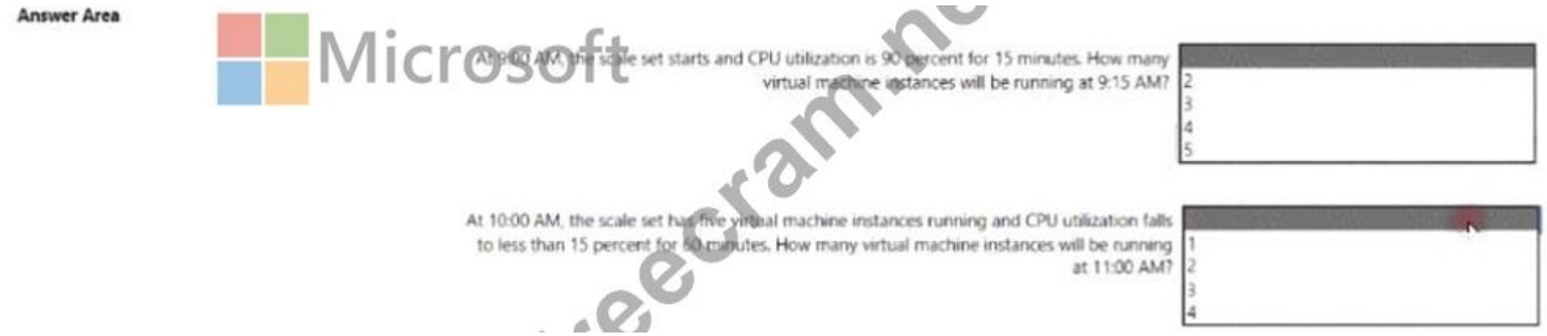

Answer:

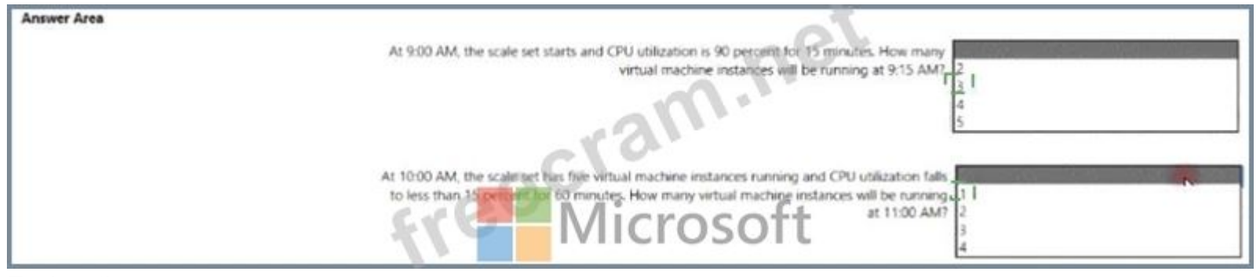

Explanation:

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: 84

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: 85
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 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 Azure AD in the Azure portal, you use the Bulk create user operation.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
This question focuses on creating guest user accounts (B2B users) in Azure Active Directory (Microsoft Entra ID) using bulk operations. The requirement is to import 500 external users from a CSV file and invite them as guest users into the tenant contoso.com.
Understanding the Goal
You must create guest accounts (B2B collaboration users) - not regular internal users.
Guest accounts in Azure AD have:
UserType = Guest
Their identity originates from an external organization (using their own email domain).
They are invited via B2B invitations.
The key operation used to accomplish this is Bulk invite users, not Bulk create users.
Why "Bulk Create Users" Does NOT Meet the Goal
The Bulk create users operation in the Azure portal is used for adding internal users to the directory (i.e., users who belong to the same Azure AD tenant).
It creates standard member accounts, not guest accounts.
It cannot send invitations to external users or create guest-type identities (UserType = Guest).
In contrast:

The Bulk invite users operation imports external user data from a CSV file and automatically sends invitations to the listed email addresses.
It creates user objects with UserType = Guest in Azure AD, meeting the scenario's goal.
Hence, the Bulk create user method fails to create guest users and does not meet the goal.
Correct Method (for Reference)
From Microsoft Entra ID # Users # Bulk operations # Bulk invite, you can:
Download a CSV template.
Fill in external users' names and email addresses.
Upload the CSV file.
Azure AD then creates guest user accounts and sends invitations to join the directory.
This is the only correct approach for mass-adding external (guest) users.
Verified Documentation Extract (Microsoft Learn)
"The Bulk invite operation in Azure AD allows administrators to invite multiple external users at once by uploading a CSV file. Bulk create operation, on the other hand, is for adding internal users to your organization's directory." (Source: Microsoft Learn - Invite external users in bulk to Microsoft Entra ID, AZ-104 Study Guide)
Final Verified Answer:
B). No
Using the Bulk create users operation adds internal member accounts, not external guest accounts. To create guest accounts, the correct action is to use the Bulk invite users operation in Azure AD.

NEW QUESTION: 86
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: 87
You have an Azure subscription that contains the resources shown in the following table.

Name	Type
storage1	Storage account
container1	Blob container
table1	Storage table

You need to perform the tasks shown in the following table.

Name	Task
Task1	Create a new storage account.
Task2	Upload an append blob to container1.
Task3	Create a file share in storage1.
Task4	Add data to table1.

Which tasks can you perform by using Azure Storage Explorer?

A. Task1, Task2, Task3, and Task4

B. Task1 and Task3 only

C. Task1, Task3, and Task4 only

D. Task1, Task2, and Task3 only

E. Task2, Task3, and Task4 only

Answer: (SHOW ANSWER)

NEW QUESTION: 88

You have an Azure subscription that contains a storage account named storage1. The storage 1 account contains a container named containet1.

You create a blob lifecycle rule named rule1.

You need to configure rule1 to automatically move blobs that were NOT updated for 45 days Irom container!

to the Cool access tier.

How should you complete the rule? To answer, select Ihe appropriate options in the answer area.

NOTE: Each correct answer is worth one point.

Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToCool": {
                                                    : 45
                "daysAfterCreationCreaterThan"
                "daysAfterLastAccessTimeGreaterThan"
                "daysAfterModificationGreaterThan"
            }
          }
        },
        "filters": {
          "blobTypes": [
                "AppendBlob"
                "Blockblob"
                "Pageblob"
          ],
          "prefixMatch": [
            "container1"
          ]
        }
      }
    }
  ]
}
```


Answer:

Answer Area

Explanation:

Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToCool": {
              "daysAfterModificationGreaterThan"  45
            }
          }
        },
        "filters": {
          "blobTypes": [
            "Blockblob"
          ],
          "prefixMatch": [
            "container1"
          ]
        }
      }
    }
  ]
}
```


Azure Blob Storage Lifecycle Management enables you to automatically transition data to cooler storage tiers, archive data, or delete data based on specified conditions related to creation, last modification, or last access time. These policies are defined in a JSON-based rule structure under the rules property.

As per Microsoft Azure Administrator Study Guide and Official Documentation ("Manage the Azure Blob storage lifecycle"), each rule must define:

Actions: What operation to perform (e.g., tierToCool, tierToArchive, delete).

Filters: Which objects (blobs) the rule applies to, typically filtered by container name, prefix, or blob type.

Conditions: When to perform the action, defined by properties such as:

" daysAfterCreationGreaterThan " - Number of days since blob creation.

" daysAfterLastAccessTimeGreaterThan " - Number of days since last read operation (used when last access tracking is enabled).

" daysAfterModificationGreaterThan " - Number of days since last modification (used for update or write operations).

According to the scenario, the requirement is to move blobs that were not updated for 45 days. The "not updated" condition refers to the last modification date, not creation or access date. Therefore, the correct property to use is:

" daysAfterModificationGreaterThan " : 45

The blob type must be specified under " filters " . Azure Lifecycle Management currently supports lifecycle actions for:

BlockBlob - the primary type for uploaded files.

AppendBlob - supported, but typically used for log data.

PageBlob - not supported for lifecycle tiering (used for virtual machine disks).

Per Microsoft Docs - "Lifecycle management policy definition", lifecycle tiering actions (tierToCool or tierToArchive) only apply to BlockBlob and AppendBlob, with BlockBlob being the standard choice for data that benefits from tier transitions between Hot, Cool, and Archive tiers.

Hence, the correct configuration is:

" tierToCool " : {

" daysAfterModificationGreaterThan " : 45

}
" blobTypes " : [
" BlockBlob "
]
This policy automatically transitions block blobs in container1 to the Cool access tier when they have not been modified for 45 days, optimizing storage cost while retaining availability.

NEW QUESTION: 89
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.

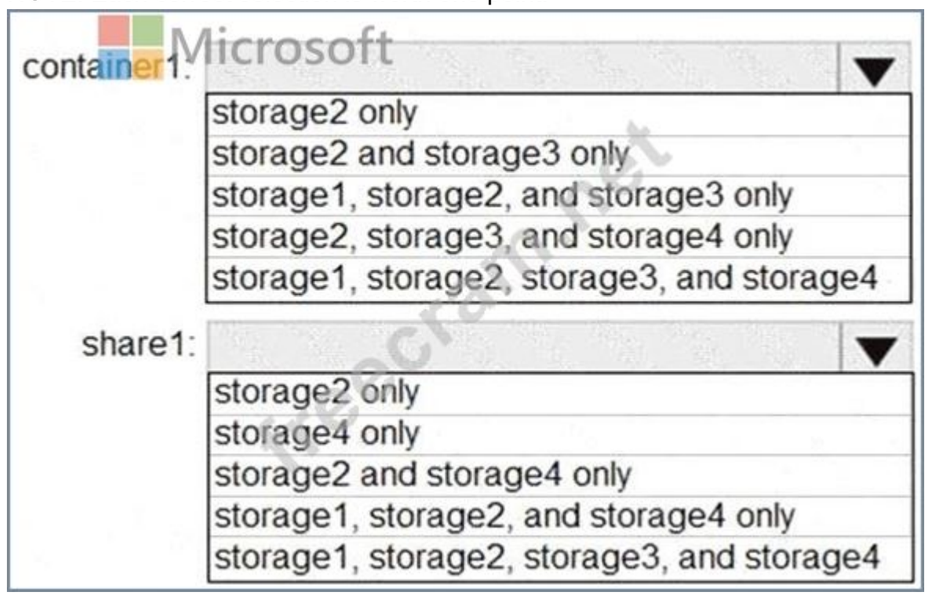

Answer:
In Azure, the type of storage account determines which services are available and which features can be configured. Let's analyze the given storage accounts in this case study:
Storage Account
Kind
Services Supported
Hierarchical Namespace
File Share Access (Identity-based)
storage1
General-purpose 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: 90

You need to configure an Azure web app named contoso.azurewebsites.net to host www.contoso.com.
What should you do first?

A. Create a CNAME record named asuid that contains the domain verification ID.
B. Create A records named www.contoso.com and asuid.contoso.com.
C. Create a TXT record named asuid that contains the domain verification ID.
D. Create a TXT record named www.contoso.com that has a value of contoso.azurewebsites.net.

Answer: A (LEAVE A REPLY)
When you configure a custom domain (like www.contoso.com) for an Azure Web App (App Service), Azure requires that the domain be verified to ensure ownership before binding it to the app.
According to Microsoft Azure official documentation ("Map a custom domain name to your Azure web app"
- Microsoft Learn):
"Before you can add a custom domain, you must verify that you own the domain name by creating a DNS record with your domain registrar. Azure uses an asuid verification record, which can be either a CNAME or TXT record, depending on your DNS provider." The first step is to create a CNAME or TXT record in your DNS zone that links your custom domain to the Azure verification ID.
The record name is asuid.contoso.com.
The value (target) is the domain verification ID shown in the Azure portal under Custom domains # Custom hostnames # Domain ownership.
Once Azure verifies the domain ownership using that record, you can add www.contoso.com as a custom hostname in the web app configuration.
Key Point:
Step 1: Verify domain ownership using a CNAME or TXT record.
Step 2: Bind the custom hostname (www.contoso.com) to your web app.
Thus, the correct and verified answer is:
A. Create a CNAME record named asuid that contains the domain verification ID.

NEW QUESTION: 91
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: 92
You have an Azure subscription named Subscription1.
You create an Azure Storage account named contosostorage, and then you create a file share named data.

Which UNC path should you include in a script that references files from the data file share? To answer, drag the appropriate values to the correct targets. Each value 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:

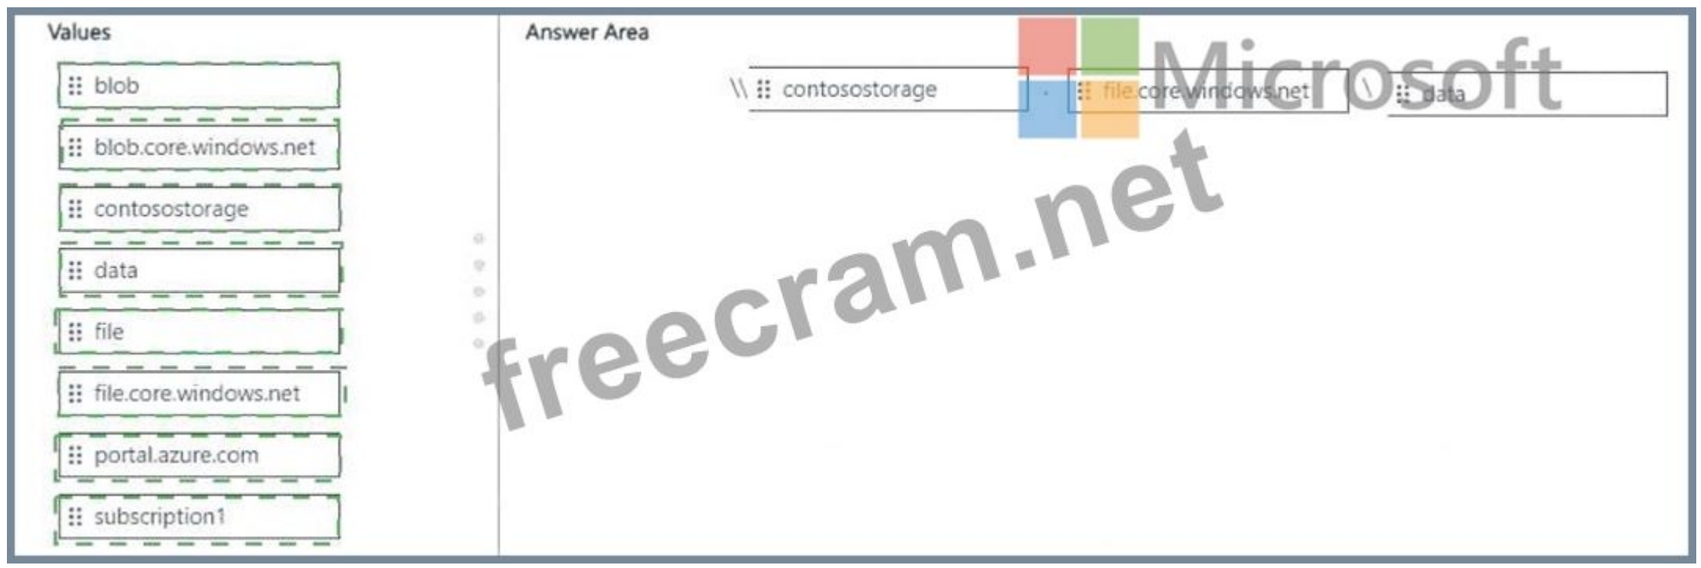

Explanation:

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

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

\\ < storage-account-name > .file.core.windows.net\ < file-share-name > In this scenario:

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

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

NEW QUESTION: 93

You plan to create the Azure web apps shown in the following Table.

Name	Runtime stack
WebApp1	.NET 6 (LTS)
WebApp2	ASP.NET V4.8
WebApp3	PHP 8.1
WebApp4	Python 3.11

What is the minimum number of App Service plans you should create for the web apps?

A. 1

B. 2

C. 3

D. 4

Answer: B (LEAVE A REPLY)

NET Core 3.0: Windows and Linux ASP .NET V4.7: Windows only PHP 7.3: Windows and Linux Ruby 2.6:

Linux only Also, you can't use Windows and Linux Apps in the same App Service Plan, because when you create a new App Service plan you have to choose the OS type. You can ' t mix Windows and Linux apps in the same App Service plan. So, you need 2 ASPs. Reference: https://docs.microsoft.com/en-us/azure/app- service/overview

NEW QUESTION: 94

You have an Azure Load Balancer named LB1.

You assign a user named User1 the roles shown in the following exhibit.

Answer:

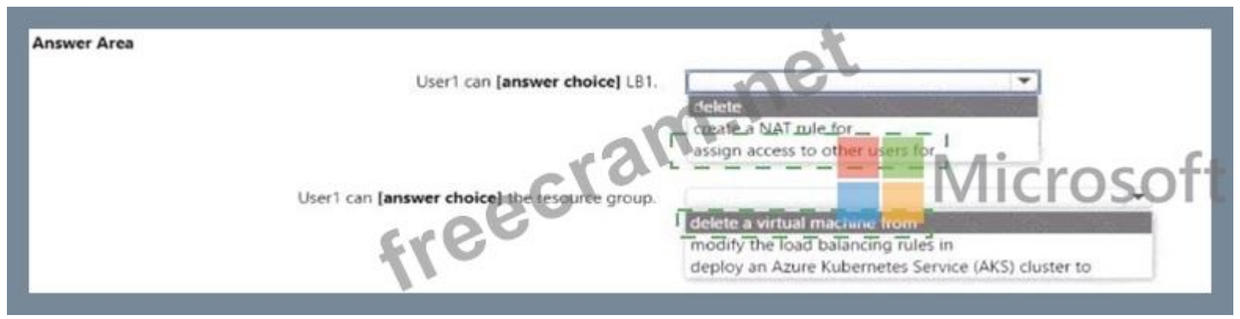

Explanation:

User Access Administrator can only assign access to other users

https://docs.microsoft.com/en-us/azure/role-based-access-control/rbac-and-directory-admin-roles Virtual Machine Contributor can Manage VMs, which includes deleting VMs too.

https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles#virtual-machine-contributor

https://docs.microsoft.com/en-us/answers/questions/350635/can-virtual-machine-contributor-create-vm.html

NEW QUESTION: 95

You need to configure WebApp1 to meet the technical requirements.

Which certificate can you use from Vault1?

A. Cert1 only

B. Cert1 or Cert2 only

C. Cert1 or Cert3 only

D. Cert3 or Cert4 only

E. Cert1, Cert2, Cert3, or Cert4

Answer: A (LEAVE A REPLY)

To meet the technical requirement - "Use TLS for WebApp1" - the web app must be configured with a certificate that is compatible with Azure App Service for HTTPS/TLS binding.

According to the Microsoft Azure Administrator documentation on App Service Certificates and Key Vault integration, the following key points determine which certificates can be used:

* Supported Certificate Format:Azure App Service supports importing certificates in PFX (PKCS #12) format, which includes both the public and private keys necessary for TLS/SSL binding.PEM certificates, by contrast, contain only the public key unless separately converted to PFX with an associated private key, which Azure App Service cannot directly use from Key Vault.

* Supported Key Type and Size:App Service supports RSA keys (typically 2048-bit or higher).Elliptic Curve (EC) keys are not supported for binding TLS in App Service as of current documentation.

* Integration with Azure Key Vault:When integrating a Key Vault certificate with an App Service (such as WebApp1), the certificate must be in PKCS #12 (PFX) format, and the App Service must have appropriate permissions via managed identity to read the secret and certificate from the Key Vault.

From the Vault1 data provided in your scenario:

Name

Content type

Key type

Key size

Cert1

PKCS #12

RSA

2048

Cert2

PKCS #12

RSA

4096

Cert3

PEM

RSA

2048

Cert4

PEM

RSA

4096

Analysis:

* Cert1 and Cert2 are PKCS #12 certificates, so both contain the private key required for TLS.

* However, only Cert1 (RSA 2048) is a Microsoft-recommended configuration for Azure Web App SSL

/TLS use.

* Cert2 has a 4096-bit RSA key. Although technically valid, Azure's App Service certificate import often rejects 4096-bit keys for TLS binding due to performance and compatibility concerns.

* Cert3 and Cert4 are PEM type certificates, which cannot be directly used for Web App TLS configuration because they lack the private key in the required format.

Therefore, according to the Azure Administrator Exam Study Guide and Microsoft official documentation, the only valid certificate that meets the requirements is:
Cert1 only
Final Verified Answer: # A. Cert1 only
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: 96

You have an Azure container registry in the Standard tier.

You need to enable the admin user account and retrieve the password for the account Which settings should you use?

A. Networking

B. Properties
C. Identity
D. Access keys
Answer: (SHOW ANSWER)
In Azure Container Registry (ACR), the admin user account provides a simple authentication mechanism using a username and password. Enabling the admin account and retrieving its credentials is done exclusively through the Access keys blade.
According to Azure Container Registry documentation:
The admin user is enabled/disabled from Access keys
The passwords (primary and secondary) are displayed and regenerated there Other blades such as Networking, Properties, or Identity do not manage admin credentials Final Answer:
D). Access keys

NEW QUESTION: 97
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: 98
You have an Azure subscription that contains the resources shown in the following table.

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

You plan to create a data collection rule named DCRI in Azure Monitor.
Which resources can you set as data sources in DCRI, and which resources can you set as destinations in DCRI? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

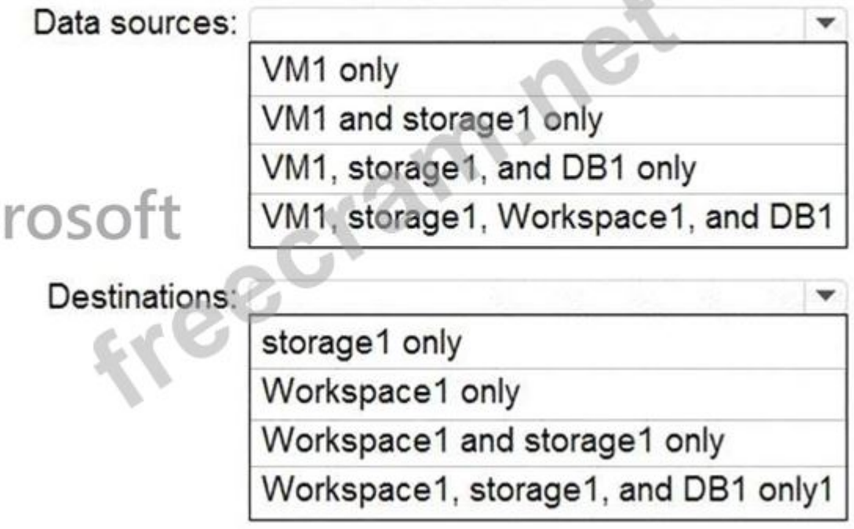

Answer:

Answer Area

Data sources:

- VM1 only
- VM1 and storage1 only
- VM1, storage1, and DB1 only
- VM1, storage1, Workspace1, and DB1

Destinations:

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

Explanation:

Data Sources: VM1 only

Destination: Workspace1 Only

NEW QUESTION: 99

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:

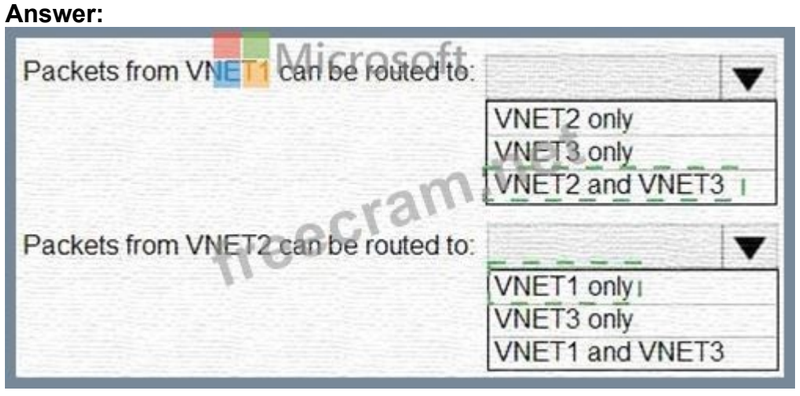

Explanation:

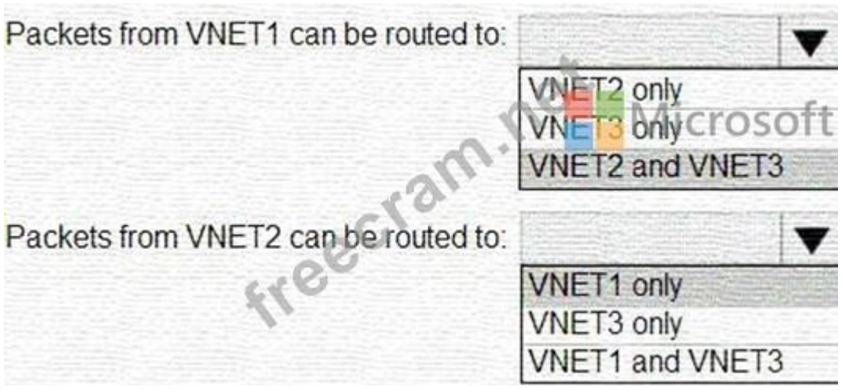

This question tests your understanding of Azure Virtual Network (VNet) Peering and its transitivity limitations.

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: 100
You have an Azure subscription that contains a storage account named storage1. The subscription is linked to a Microsoft Entra tenant named contoso.com that syncs with an on-premises Active Directory domain. The domain contains the security principals shown in the following table.

Name	Type
User1	User
Computer1	Computer

In the Microsoft Entra tenant you create a user named User2.
The storage1 account contains a file share named share! and has the following configurations.
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:

NEW QUESTION: 101

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:

Answer Area

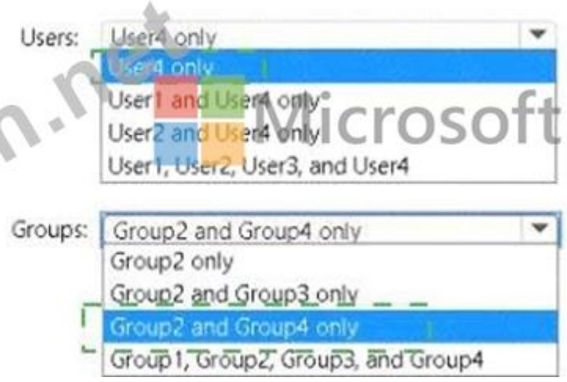

Explanation:

Answer Area

NEW QUESTION: 102

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. the New-AzConfigurationAssignment Cmdlet

B. Azure Application Insights

C. the Publish-ArVMDscConfiguration cmdlet

D. Azure Custom Script Extension

Answer: (SHOW ANSWER)

NEW QUESTION: 103

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

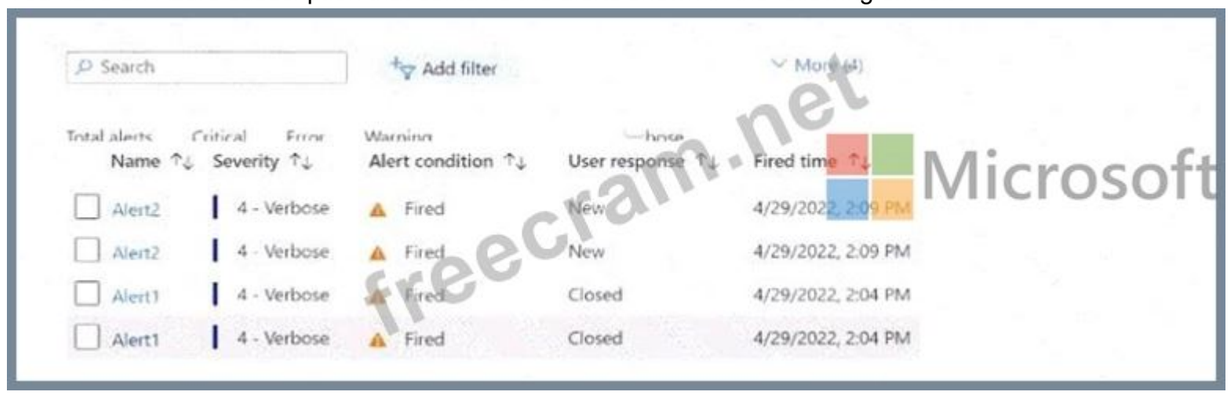

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE Each correct selection is worth one point.

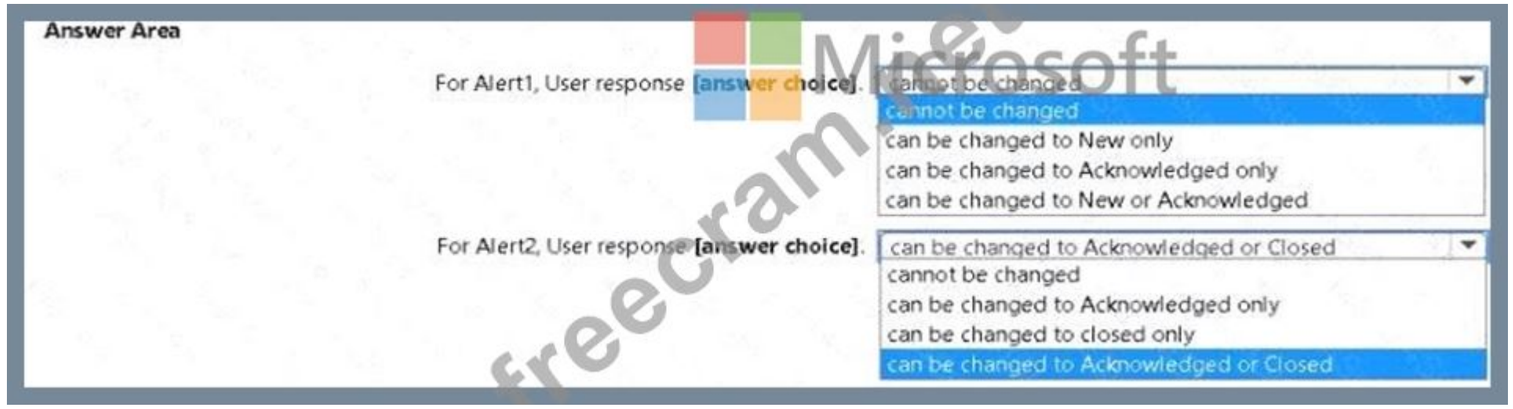

Answer:

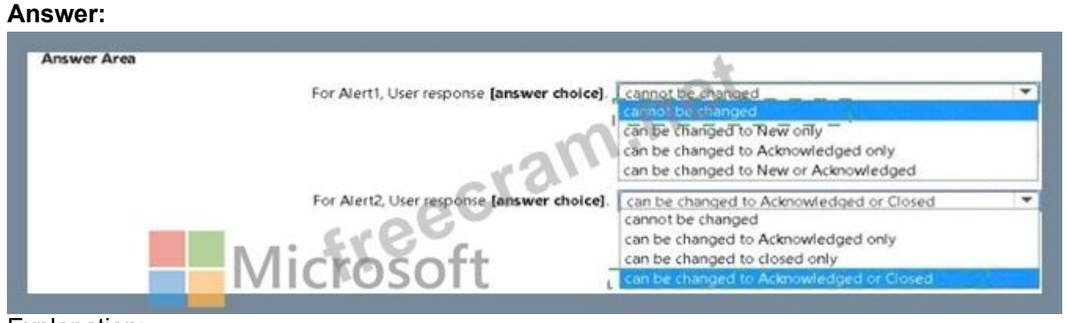

Explanation:

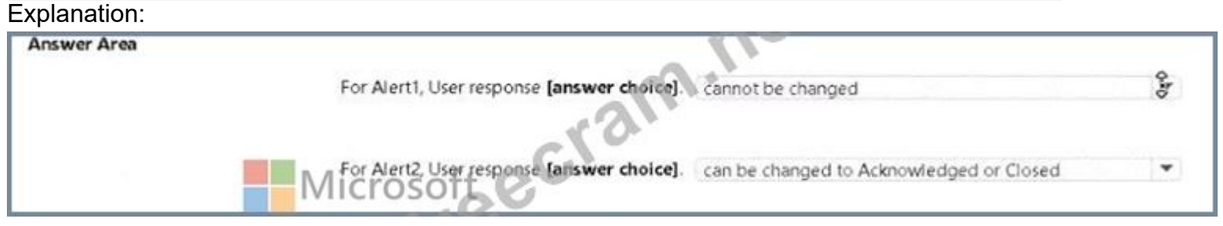

In Azure Monitor, when an alert rule triggers, an alert instance is created. Each alert has a state and a user response status that helps track its lifecycle. The alert state (such as Fired or Resolved) reflects the current condition of the monitored resource, while the user response (such as New, Acknowledged, or Closed) indicates the operator's handling of the alert.
According to the Microsoft Azure Administrator study guide and official Azure Monitor documentation, once an alert has been closed, it enters a terminal state where no further status modifications are permitted - this is by design to preserve operational tracking integrity. Therefore, Alert1, whose user response is "Closed," cannot be changed.
For Alert2, which is in a "New" state, the response can be updated as part of the triage workflow.
Administrators can manually change the alert's user response to Acknowledged (indicating investigation is underway) or to Closed (indicating the issue is resolved).
This structure aligns with best practices outlined in the Azure Monitor alert lifecycle management documentation, ensuring consistent tracking, accountability, and auditability across alert operations.
Final Verified Answers:
* Alert1 (User response): cannot be changed
* Alert2 (User response): can be changed to Acknowledged or Closed #

NEW QUESTION: 104
You have an Azure Storage account named storage1.
You plan to use AzCopy to copy data to storage1.
You need to identify the storage services in storage1 to which you can copy the data.
What should you identify?
A. blob, file, table, and queue
B. blob and file only
C. file and table only
D. file only
E. blob, table, and queue only

Answer: (SHOW ANSWER)
AzCopy is a command-line utility provided by Microsoft Azure for transferring data to and from Azure Storage accounts with high performance and reliability. It is particularly designed to support bulk data migration and synchronization for Azure storage services.
According to the Microsoft Azure Administrator documentation (AzCopy v10 and later), the tool supports copying data to and from only two Azure storage services:
* Azure Blob storage - including block blobs, append blobs, and page blobs.
* Azure Files (File shares) - including directories and individual files.
The Table and Queue storage services are not supported by AzCopy, as these services handle structured and message-based data rather than binary or file data.
From the official Azure Administrator Study Guide and Microsoft Docs:
"AzCopy is a command-line tool that moves data into and out of Azure Storage accounts. You can use AzCopy with Blob storage or Azure Files. AzCopy doesn't support copying data to or from Table storage or Queue storage." (Source: Microsoft Learn - "Transfer data with AzCopy and Blob storage" and "Transfer data with AzCopy and Azure Files")

NEW QUESTION: 105
You have a Microsoft Entra tenant.
You need to modify the Default user role permissions settings for the tenant. The solution must meet the following requirements:
* Standard users must be prevented from creating new service principals.
* Standard users must only be able to use PowerShell or Microsoft Graph to manage their own Azure resources.
Which two settings should you modify? To answer, select the appropriate settings in the answer area.
NOTE: Each correct answer is worth one point.

Answer:

Explanation:

1. Users can register applications # No

2. Restrict access to Microsoft Entra ID administration portal # Yes

In Microsoft Entra ID (formerly Azure Active Directory), the Default User Role Permissions setting determines what standard (non-admin) users can do in the directory by default. To meet the two stated

requirements - Prevent standard users from creating new service principals (applications), and Restrict them to managing only their own Azure resources using PowerShell or Microsoft Graph - you must modify the following two settings:
1. Users can register applications # No
When this setting is enabled, any user in the directory can register applications, which automatically creates service principals in Microsoft Entra ID. These service principals can be used to grant access to resources, effectively creating new security identities.
According to the official Microsoft Entra Administrator documentation:
"If you want to prevent non-administrators from creating app registrations or service principals in your organization, disable the 'Users can register applications' setting under Default User Role Permissions." By setting Users can register applications = No, you prevent standard users from creating new app registrations or service principals, fulfilling the first requirement.
2. Restrict access to Microsoft Entra ID administration portal # Yes
This setting determines whether standard users can access the Microsoft Entra admin center (portal.azure.
com). When you enable Restrict access to Microsoft Entra ID administration portal, standard users can no longer access or modify directory-wide settings through the portal.
However, they can still manage their own Azure resources via PowerShell or Microsoft Graph (within their permissions).
As documented in Microsoft Learn (Manage default user permissions in Microsoft Entra ID):
"To ensure that users can only manage their own objects and resources, but not the directory configuration, enable 'Restrict access to Microsoft Entra ID administration portal.' This restriction does not prevent the use of PowerShell or Graph API for personal resources." This meets the second requirement that standard users can only use automation tools to manage their own Azure resources.
Summary of Correct Settings:
Setting
Correct Value
Reason
Users can register applications
No
Prevents creation of service principals by standard users.
Restrict access to Microsoft Entra ID administration portal
Yes
Ensures users only manage their own resources using PowerShell or Graph.
Final Verified Answer:
1. Users can register applications # No
2. Restrict access to Microsoft Entra ID administration portal # Yes

NEW QUESTION: 106
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 Encryption Scope Contributor Role to User1.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In Azure, storage account keys are managed through Azure Role-Based Access Control (RBAC) permissions.
The ability to list and regenerate storage account keys requires access to specific management operations on the storage account resource, not encryption scopes.
According to Microsoft Azure Administrator documentation, the permissions required to list and regenerate storage account keys are:
* Microsoft.Storage/storageAccounts/listkeys/action

* Microsoft.Storage/storageAccounts/regeneratekey/action
Only specific built-in roles include these actions, such as:
* Owner
* Contributor
* Storage Account Contributor
However, the role mentioned in the question - Storage Account Encryption Scope Contributor - is designed for a different purpose.
From Microsoft's official documentation for Storage Account Encryption Scope Contributor:
"Grants permissions to manage encryption scopes within a storage account. Does not grant access to manage storage account keys or data." This means that assigning the Storage Account Encryption Scope Contributor role gives the user permissions to configure encryption scopes (for example, enabling or modifying encryption at the blob container level), but not to list or regenerate access keys.
Therefore, the solution fails to meet the requirement of enabling User1 to list and regenerate keys.
To achieve the goal, the correct role would be:
Storage Account Contributor
This role explicitly allows managing all storage account properties, including key management.
Key Point (Microsoft Official Documentation Summary):
* Storage Account Encryption Scope Contributor # Manage encryption scopes only.
* Storage Account Contributor # Manage storage account settings, including listing and regenerating keys.

NEW QUESTION: 107
You have an Azure subscription that contains a storage account named account1.
You plan to upload the disk files of a virtual machine to account! from your on-premises network. The on- premises network uses a public IP address space of 131.107.1.0/24.
You plan to use the disk files to provision an Azure virtual machine named VM1. VM1 will be attached to a virtual network named VNet1. VNet1 uses an IP address space of 192.168.0.0/24.
You need to configure account1 to meet the following requirements:
* Ensure that you can upload the disk files to account1.
* Ensure that you can attach the disks to VM1.
* Prevent all other access to account1.
Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.
A. From the Networking blade of account1, select Selected networks
B. From the Service endpoints blade of VNet1, add a service endpoint.
C. From the Networking blade of account11, add the 131.107.1.0/24 IP address range.
D. From the Networking blade of account1. select Allow trusted Microsoft services to access this storage account
E. From the Networking blade of account1, add VNet1.
Answer: (SHOW ANSWER)
To restrict access to account1, you need to enable the firewall and virtual network settings on the storage account. This allows you to specify which networks can access the storage account. By selecting Selected networks, you can block all access from the public internet and only allow access from the specified networks. By adding VNet1, you can allow access from the virtual network that contains VM1. You do not need to add the on-premises IP address range or enable the service endpoint option, as these are not required for uploading the disk files to the storage account. You do not need to allow trusted Microsoft

services, as this is not relevant for the scenario. Then, References: [Configure Azure Storage firewalls and virtual networks] [Upload a generalized VHD to Azure]

NEW QUESTION: 108
You have an Azure policy as shown in the following exhibit.

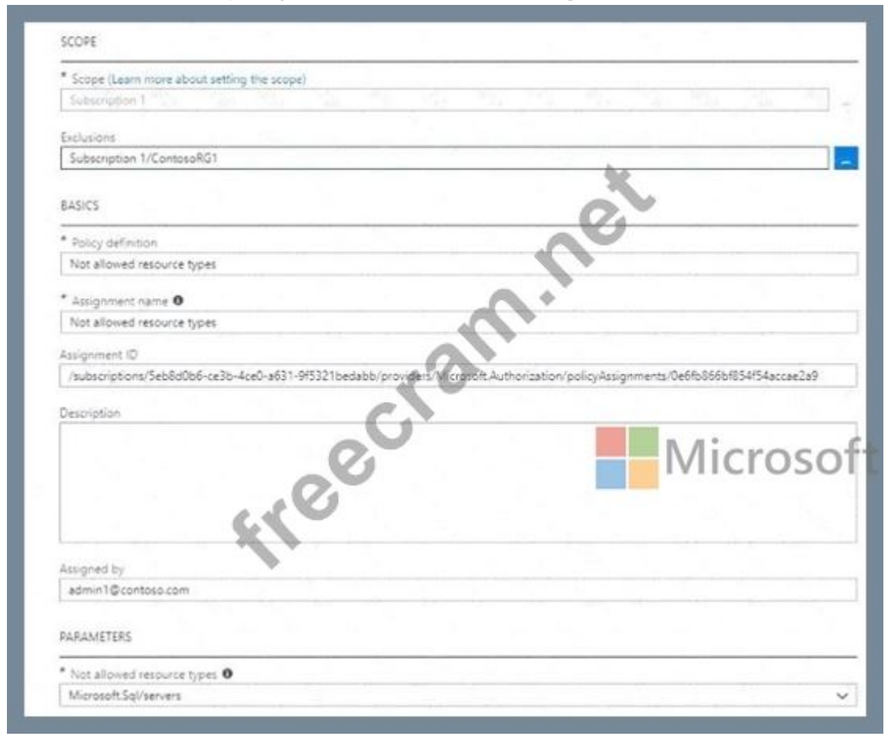

What is the effect of the policy?
A. You can create Azure SQL servers in ContosoRG1 only.
B. You can create Azure SQL servers in any resource group within Subscnption1.
C. You are prevented from creating Azure SQL Servers in ContosoRG1 only.
D. You are prevented from creating Azure SQL servers anywhere in Subscnption1.
Answer: (SHOW ANSWER)

NEW QUESTION: 109
You have several Azure virtual machines on a virtual network named VNet1.
You configure an Azure Storage account as 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 Area

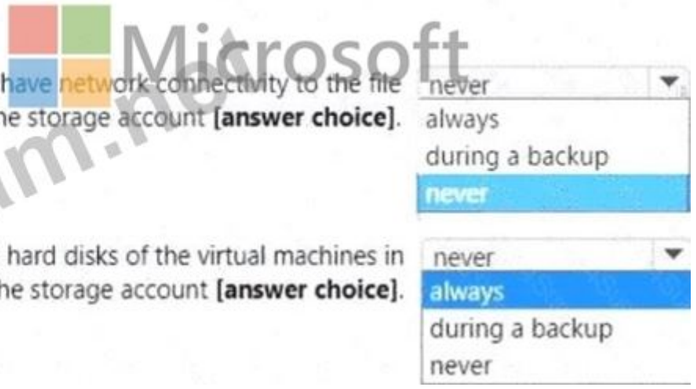

Answer:

Answer Area

Explanation:

Never

Never

NEW QUESTION: 110

You plan to deploy the following Azure Resource Manager (ARM) template.

```
{
  "$schema":"https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
  "contentVersion":"1.0.0.0",
  "parameters":{},
  "variables":{
    "vnetId":"[resourceId('Microsoft.Network/virtualNetworks/', 'VNET1')]",
    "lbId":"[resourceId('Microsoft.Network/loadBalancers/', 'LB1')]",
    "sku":"Standard",
    "netname":"APP1"
  },
  "resources":[
    {
      "apiVersion":"2017-08-01",
      "type":"Microsoft.Network/loadBalancers",
      "name":"LB1",
      "location":"EastUS",
      "sku":{
        "name":"[variables('sku')]"
      },
      "properties":{
        "frontendIPConfigurations":[
          {
            "name":"[variables('netname')]",
  }
],
"loadBalancingRules":[
    properties :{
      "frontendIPConfiguration":{
        "id":"[concat(variables('lbId'),'/frontendIPConfigurations/', variables('netname'))]"
      },
      "backendAddressPool":{
        "id":"[concat(variables('lbId'), '/backendAddressPools/',variables('netname'),'-Servers')]"
      },
      "probe":{
        "id":"[concat(variables('lbId'), '/probes/probe')]"
          },
          "backendPort":8080,
          "protocol":"Tcp",
          "frontendPort":80,
          "enableFloatingIP":false,
          "idleTimeoutInMinutes":4,
          "loadDistribution":"SourceIPProtocol"
        }
      }
    ],
    "probes":[
      {
        "name":"probe",
        "properties":{
          "protocol":"Tcp",
          "port":8080,
          "intervalInSeconds":15,
          "numberOfProbes":2
        }
      }
```

For each of the following statements, select Yes . Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
LB1 will be connected to a subnet named VNET1/netname.		
LB1 can be deployed only to the resource group that contains VNET1.		
The value of the sku variable can be provided as a parameter when the template is deployed.		

Answer:

Explanation:

LB1 will be connected to a subnet named LB1 in VNET1. Yes, this is correct. The template specifies that the load balancer resource named LB1 has a property called frontendIPConfigurations, which defines the subnet where the load balancer is located. The value of this property is a reference to the resource ID of the subnet named LB1 in VNET1. You can see this reference in line 38 of the template1.

LB1 can be deployed only to the resource group that contains VNET1. No, this is not correct. The template does not specify a resource group for the load balancer resource, which means it can be deployed to any resource group in the same subscription as VNET1. However, if you want to deploy the load balancer to a specific resource group, you can add a property called resourceGroup to the reference of the subnet in line

382.

The value of the sku variable can be provided as a parameter when the template is deployed. No, this is not correct. The template defines the sku variable as a constant value of "Standard" in line 9. This means that the value cannot be changed or overridden by a parameter when the template is deployed. If you want to make the sku value configurable, you need to change the variable definition to a parameter definition, and use the parameter reference instead of the variable reference in line 363.

NEW QUESTION: 111

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: 112

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

Name	Type
ManagementGroup1	Management group
RG1	Resource group
9c8bc1cd-7655-4c66-b3ea-a8ee101d8f75	Subscription ID
Tag1	Tag

In Azure Cloud Shell, you need to create a virtual machine by using an Azure Resource Manager (ARM) template.

How should you complete the command? To answer, select the appropriate options in the answer area, NOTE: Each correct selection is worth one point.

Answer:

Explanation:

Reference:
https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-6.6.0

NEW QUESTION: 113

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: 114
You have an Azure subscription that contains the vaults shown in the following table.

Name	Type
Recovery1	Recovery Services vault
Backup1	Azure Backup vault

You deploy the virtual machines shown in the following table.

Name	Type	In vault
Policy1	Standard	Recovery1
Policy2	Enhanced	Recovery1
Policy3	*Not applicable*	Backup1

Each of the following statements, select Yes if the statement is true. Otherwise, select No NOTE: Each cored selection it worth one point.
Answer Area

Statements	Yes	No
VM1 can be backed up by using Policy1.	○	○
VM2 can be backed up by using Policy3.	○	○
VM2 can be backed up by using Policy2.	○	○

Answer:
Answer Area

Statements	Yes	No
VM1 can be backed up by using Policy1.	○	○
VM2 can be backed up by using Policy3.	○	○
VM2 can be backed up by using Policy2.	○	○

Explanation:

Azure provides two types of vaults for backup management - Recovery Services vaults and Backup vaults.
Each vault type supports different backup management capabilities and policies:
Recovery Services vault
This is the modern, recommended vault for managing backup and recovery across a wide range of Azure resources including Azure Virtual Machines, Azure Files, and on-premises workloads.
Within a Recovery Services vault, you can configure Standard and Enhanced backup policies.
Standard Policy supports backup schedules and retention policies using Azure Backup's classic model.
Enhanced Policy introduces granular backup capabilities and supports backups more frequently, using the modernized Azure Backup architecture.
As per Azure documentation:
"The Recovery Services vault supports both Standard and Enhanced policies for backing up Azure VMs.
Enhanced policies offer improved flexibility and reliability and can coexist with Standard policies in the same vault." (Source: Microsoft Azure Backup Architecture Guide and Recovery Services Vault Overview - Azure Documentation, AZ-104 Study Guide) Azure Backup vault Backup vaults are specific to Azure Backup (Modern) workloads such as Azure Resource Manager-based backups for Azure Files and Azure Disks. However, these cannot directly back up Azure VMs.
Backup vaults were introduced to support "Azure Backup Center" and centralized management but are not interchangeable with Recovery Services vaults for VM backups. Policies created in an Azure Backup vault (e.
g., Policy3) cannot be applied to VMs managed under a Recovery Services vault and vice versa.
Therefore, in this scenario:
VM1 is backed up in Recovery1 (a Recovery Services vault). Hence, it can use both Policy1 (Standard) and Policy2 (Enhanced).
VM2 is also managed under Recovery1, so it can also use Policy2.
Policy3, located in Backup1 (an Azure Backup vault), cannot be used for VM backups in Recovery1.
Thus, the correct answers are:
VM1 # Policy1 # # Yes
VM2 # Policy3 # # No
VM2 # Policy2 # # Yes

NEW QUESTION: 115
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: 116

You have an Azure subscription that has the Azure container registries shown in the following table.

Name	Service tier
ContReg1	Premium
ContReg2	Standard
ContReg3	Basic

You plan to use ACR Tasks and configure endpoint connections.

Answer:

Explanation:

ACR Tasks: ContReg1 and ContReg2 only

Private endpoints: ContReg1 only

Azure Container Registry (ACR) provides multiple service tiers (SKUs)-Basic, Standard, and Premium- each offering different capabilities for scalability, performance, and features. These tiers determine what functionalities are available, including ACR Tasks and Private Link (Private Endpoints) integration.

Here's how each tier differs according to Microsoft Azure Administrator documentation (AZ-104 Exam Study Guide and Azure Docs):

Feature

Basic

Standard

Premium

ACR Tasks (Build automation, image building, and updating)

Supported

Supported

Supported

Private Link / Private Endpoints (Private access via Azure backbone)

Not Supported

Not Supported

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

NEW QUESTION: 117
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.

Answer:

Explanation:

NYN

NEW QUESTION: 118

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

NEW QUESTION: 119
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 Network Watcher, you create a connection monitor.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
In this scenario, you need to inspect all network traffic between VM1 and VM2 in VNet1 for a period of time.
The proposed solution uses Azure Network Watcher - Connection Monitor. However, this solution does not meet the goal, because Connection Monitor is designed to test connectivity and monitor latency, packet loss, and reachability between two endpoints-not to capture or inspect the actual contents or packets of the traffic.
According to Microsoft Azure Administrator Study Guide and Azure Network Watcher official documentation, Azure provides different tools for different purposes:
Connection Monitor: Verifies that a connection exists between two VMs or endpoints and monitors metrics such as latency, availability, and packet loss. It does not capture network packets or provide detailed traffic inspection.
Network Watcher Packet Capture: Captures actual network packets entering or leaving a virtual machine. It is the correct tool to use when you need to inspect all network traffic between VMs. Packet capture can be configured to run for a specific time (for example, three hours) and stored in Azure Storage for later analysis.
From the Microsoft Learn: " Implement and Manage Network Watcher " documentation, the correct approach is:
"Use Network Watcher Packet Capture to capture network traffic to and from a virtual machine. Packet capture helps in diagnostics by collecting network traffic over a specified time frame or under certain conditions." Therefore, to meet the goal of inspecting all traffic between VM1 and VM2, you must use Packet Capture, not Connection Monitor.

NEW QUESTION: 120
You have an Azure subscription named Subscription1 that contains a resource group named RG1.
In RG1, you create an internal load balancer named LB1 and a public load balancer named LB2.
You need to ensure that an administrator named Admin1 can manage LB1 and LB2. The solution must follow the principle of least privilege.
Which role should you assign to Admin1 for each task? 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 RBAC (Role-Based Access Control) and the principle of least privilege when delegating permissions to manage specific load balancers.

Scenario Summary

You have:

Resource group: RG1

Two Load Balancers:

LB1 (Internal)

LB2 (Public)

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

Add a backend pool to LB1

Add a health probe to LB2

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

Understanding the Role Requirements

1## Adding a Backend Pool to a Load Balancer
To add or modify a backend pool, you need permissions to:
Update the load balancer's configuration
Modify the associated NIC or VM backend association
The Network Contributor role includes these permissions.
Microsoft Learn - Network Contributor role permissions:
"Grants full access to manage network resources, including virtual networks, load balancers, network interfaces, and public IP addresses. Does not grant access to manage virtual machines or storage accounts."
This means Network Contributor on LB1 (the load balancer resource itself) is sufficient to:
Add or remove backend pools
Configure load-balancing rules
Update frontend or backend associations
Correct Role: Network Contributor on LB1
No need for Contributor or Owner at the resource group level because that would grant more privileges than necessary (violates least privilege principle).
2## Adding a Health Probe to a Load Balancer
A health probe is a property of the load balancer resource itself.
To add or modify a health probe, you only need permissions to update the Load Balancer configuration.
Again, the Network Contributor role includes the required permissions to create or modify health probes.
Correct Role: Network Contributor on LB2
No additional access to RG1 or global resources is required.
Why Not Other Roles?
Role
Description
Too much / Too little
Contributor on RG1
Full access to all resources in the group
Too broad
Owner on LB1/LB2
Includes delete and permission management rights
Overprivileged
Network Contributor on RG1
Manage all network resources in RG1
Broader than needed
Network Contributor on LB1/LB2
Manage only the specified load balancer
Least privilege, correct
Final Verified Answers
Task
Role Assignment
To add a backend pool to LB1
Network Contributor on LB1
To add a health probe to LB2

Network Contributor on LB2
Microsoft Official Documentation Extracts (Azure RBAC & Load Balancer):
"The Network Contributor role can manage network resources but not grant access to others."
"To configure backend pools, health probes, and load balancing rules, assign the Network Contributor role on the load balancer resource itself."
"Follow the principle of least privilege by assigning resource-level roles instead of group- or subscription- level roles."
Final Verified Answer Summary:
Backend pool (LB1): Network Contributor on LB1
Health probe (LB2): Network Contributor on LB2

NEW QUESTION: 121
You have an Azure subscription that contains a storage account named contoso?02 3. The Contoso 2023 storage account contains the resources shown in the following table.
The Contoso 2023 storage account is configured as shown in the following exhibit.

You have a Microsoft Entra tenant that contains the users shown in the following table.

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

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
No
In the given scenario, the storage account "contoso2023" has Allow storage account key access set to Disabled. This means that access keys and account-level Shared Access Signatures (SAS) cannot be used to authorize requests to the storage account. Only Azure AD authentication and user delegation SAS tokens are supported when key-based access is disabled.
Let's analyze each user according to Azure's documented storage access behavior:
* User1 - User delegation SAS:
* A user delegation SAS is based on Microsoft Entra (Azure AD) credentials and requires blob service (not file service).
* Since the user delegation SAS can only access Blob storage, and the question specifies cont1 (Blob container), this would normally allow access if Azure AD permissions exist. However, if no RBAC permissions are granted, access is denied by default.
* User2 - Service SAS:
* A service SAS uses a storage account key to sign the token.
* When Allow storage account key access is disabled, service SAS becomes invalid, because it relies on the account key.
* Therefore, User2 cannot access blob or file data using a service SAS in this configuration.
* User3 - Account SAS:
* An account SAS is also generated using the storage account key.
* With key access disabled, account SAS tokens are blocked, and any operation using them fails with an authorization error.
* Thus, User3 cannot access the content either in the blob container or the file share.
According to Microsoft Azure Administrator Documentation:
"When you disable storage account key access, requests that use the account access keys for authorization fail. Shared access signatures (SAS) that are signed with an account key or a service SAS become invalid.
Only user delegation SAS tokens, which use Azure AD credentials, are supported." This aligns with the behavior described in Azure Storage security documentation and AZ-104 study guides.

NEW QUESTION: 122
You have an Azure App Services web app named App1.
You plan to deploy App1 by using Web Deploy.
You need to ensure that the developers of App1 can use their Azure Active Directory (Azure AD) credentials to deploy content to App1. The solution must use the principle of least privilege.
What should you do?
A. Configure app-level credentials for FTPS.
B. Assign The Website Contributor role to the developers.
C. Assign the Owner role to the developers.
D. Configure user-level credentials for FTPS.
Answer: (SHOW ANSWER)
When deploying to Azure App Service, several authentication methods exist:
* App-level FTPS credentials: Common to all users, not tied to Azure AD.
* User-level FTPS credentials: Separate per user, but still not Azure AD integrated.
* Azure Role-Based Access Control (RBAC): Allows assigning permissions through Azure AD roles.
To allow developers to deploy code to an App Service (via Web Deploy, Git, or Visual Studio), the least privilege role that grants this capability is the Website Contributor role.
Role capabilities:
* The Website Contributor role allows a user to manage web apps (including deployment, restart, and configuration changes), but not delete or grant access.
* It uses Azure AD authentication through RBAC, satisfying the requirement that developers use their Azure AD credentials.
Other options:
* Owner: Grants excessive permissions (manage access and delete resources) #
* App-level/User-level FTPS credentials: Use username/password, not Azure AD identities # Therefore, the correct and least-privileged role for developers to deploy via Web Deploy is Website Contributor.
# Final Verified Answer: B. Assign the Website Contributor role to the developers

NEW QUESTION: 123
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?
B. Swap the slots
A. Redeploy App1
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: 124
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: 125
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: 126

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.

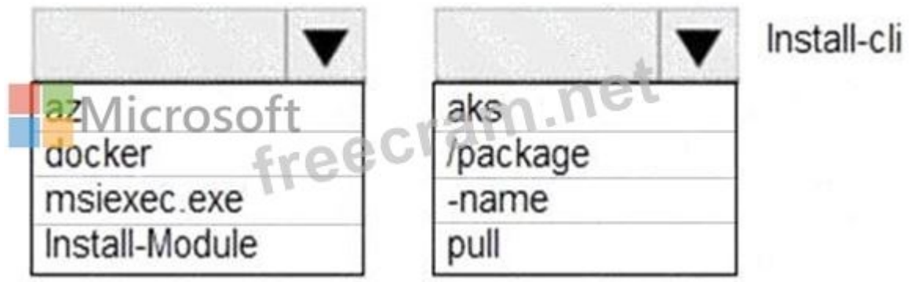

Answer:

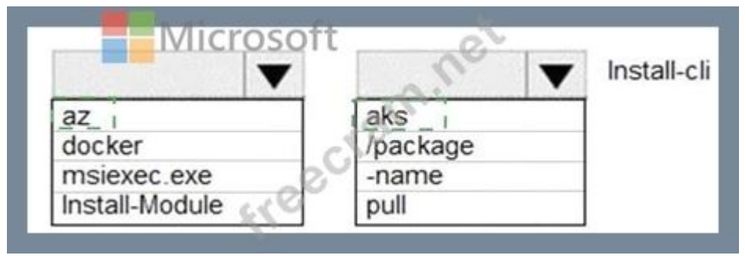

Explanation:

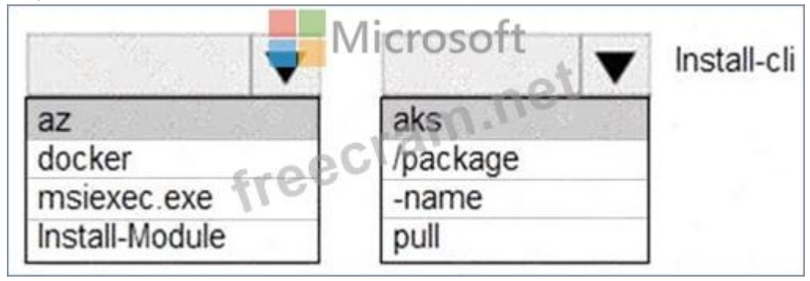

To manage and interact with an Azure Kubernetes Service (AKS) cluster from your local computer, you must have the kubectl client installed. The kubectl command-line tool lets you run commands against Kubernetes clusters to deploy applications, inspect and manage cluster resources, and view logs.

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: 127
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:

```
"assignableScopes": [
    [ "/"
      "/subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e"
      "/subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e/resourceGroups" ]
 ],
 "permissions": [
    {
      "actions": [
      "*"
      ],
      "additionalProperties" : {},
      "dataActions": [],
      "notActions" : [
        [ "Microsoft.Authorization/*"
          "Microsft.Resources/*"
          "Microsoft.Security/*" ]
    ],
    "notDataActions": []
    }
]
```


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

```
"assignableScopes": [
    [ "/"
      "/subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e"
      "/subscriptions/c276fc76-9cd4-44c9-99a7-4fd71546436e/resourceGroups" ]
```

```
],
"permissions": [
    {
      "actions": [
      "*"
      ],
      "additionalProperties" : {},
      "dataActions": [],
      "notActions" : [
```

- "Microsoft.Authorization/*"
- "Microsft.Resources/*"
- "Microsoft.Security/*"

```
      ],
      "notDataActions": []
    }
],
```


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: 128

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.

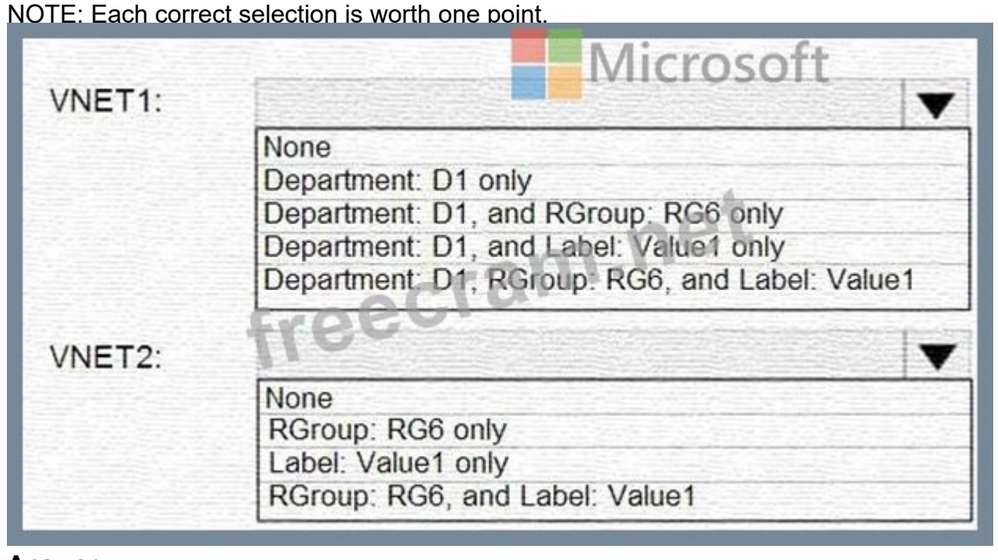

Answer:

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: 129

You have an Azure subscription.

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

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

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

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

Replication options and their redundancy levels:

Replication Type

Number of Copies

Location of Copies

Locally-redundant storage (LRS)

3

Same datacenter (single region)

Zone-redundant storage (ZRS)

3

Across availability zones in same region

Geo-redundant storage (GRS)

6

3 copies in primary region + 3 in paired region

Read-access geo-redundant storage (RA-GRS)

6

Same as GRS, plus read access to secondary region

Therefore, with LRS, the minimum number of copies is 3.

2. Reducing Cost of Infrequently Accessed Data

Azure Storage provides access tiers designed for different usage patterns:

Tier

Description

Typical Use Case

Hot

Highest storage cost, lowest access cost

Frequently accessed data

Cool
Lower storage cost, higher access cost
Infrequently accessed data
Archive
Lowest storage cost, highest retrieval cost
Long-term, rarely accessed data
As per Microsoft Learn - Manage access tiers in Azure Blob Storage:
"To reduce costs for data that is infrequently accessed, you can move blobs from the Hot tier to the Cool or Archive access tier." Hence, to minimize the cost of infrequently accessed data, you should modify the Access tier (default) setting from Hot to Cool or Archive.
Official Microsoft Extract:
From Microsoft Learn - Redundancy in Azure Storage:
"With LRS, three copies of your data exist in a single datacenter."
"To optimize storage costs for infrequently accessed data, set the access tier to Cool or Archive."
Final Verified Answers:
Statement
Correct Answer
The minimum number of copies of the storage account will be:
3
To reduce the cost of infrequently accessed data in the storage account, you must modify the:

NEW QUESTION: 130
You have an Azure subscription that contains the virtual networks shown in the following table.

Name	Location	Cloud type
VNet1	East US	Azure Government
VNet2	West US 2	Public
VNet3	China East	Azure China

You have the peering options shown in the following exhibit.

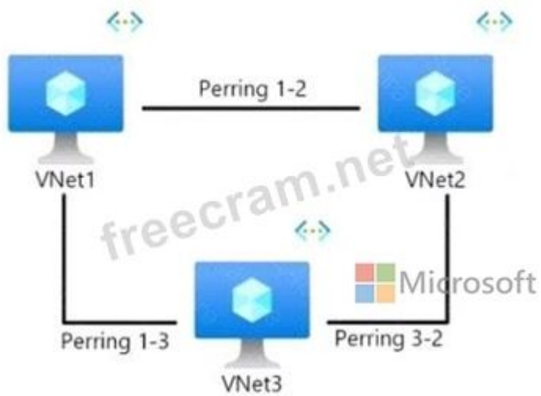

You need to design a communication strategy for the resources on the virtual networks.
For each of the following statements, select Yes if the statement is true. Otherwise, select No

Answer:

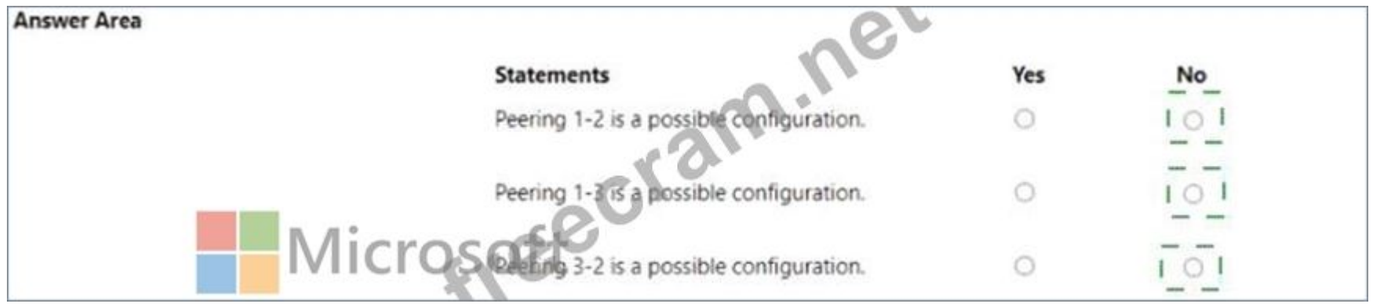

Explanation:

Azure Virtual Network (VNet) peering allows two virtual networks to communicate privately across the Azure backbone network. According to Microsoft Azure Administrator documentation, VNet peering can be of two main types:

Regional VNet Peering: within the same Azure region and cloud type.

Global VNet Peering: across Azure regions, but within the same Azure cloud.

Microsoft Azure operates multiple sovereign clouds, which are logically and physically isolated:

Azure Global/Public Cloud (used by most commercial customers).

Azure Government Cloud (for U.S. government agencies and contractors, physically isolated from the public cloud).

Azure China Cloud (operated by 21Vianet, completely independent of Microsoft's global Azure).

Per Microsoft Learn - Virtual Network Peering Documentation, the following restrictions apply:

"VNet peering cannot be established between Azure sovereign clouds (such as between Azure Government, Azure China, and Azure Public regions). Peering is only supported within the same cloud type, whether within a region or across regions." Based on the table:

VNet1 (East US, Azure Government)

VNet2 (West US 2, Public)

VNet3 (China East, Azure China)

Analysis of each peering configuration:

Peering 1-2 (Azure Government # Public Cloud): # Not supported.

Azure Government and Azure Public are separate environments - no peering across these cloud boundaries.

Peering 1-3 (Azure Government # Azure China): # Not supported.
These are two isolated sovereign clouds with independent infrastructure.
Peering 3-2 (Azure China # Public Cloud): # Not supported.
Azure China is operated by 21Vianet and cannot connect directly with the Public Azure cloud.
Therefore, no cross-cloud peering is supported among these networks.
If all VNets were within Public Azure (e.g., East US, West US 2, and North Europe), global VNet peering would be allowed. However, in this scenario, each VNet resides in a different cloud type, so no peering is possible between them.

NEW QUESTION: 131
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

Add a subnet to VNet1: User1 and User3 only
- User1 only
- User3 only
- **User1 and User3 only**
- User2 and User3 only
- User1, User2, and User3

Assign a user the Reader role to VNet1: User1 only
- **User1 only**
- User2 only
- User3 only
- User1 and User2 only
- User2 and User3 only
- User1, User2, and User3

Answer:

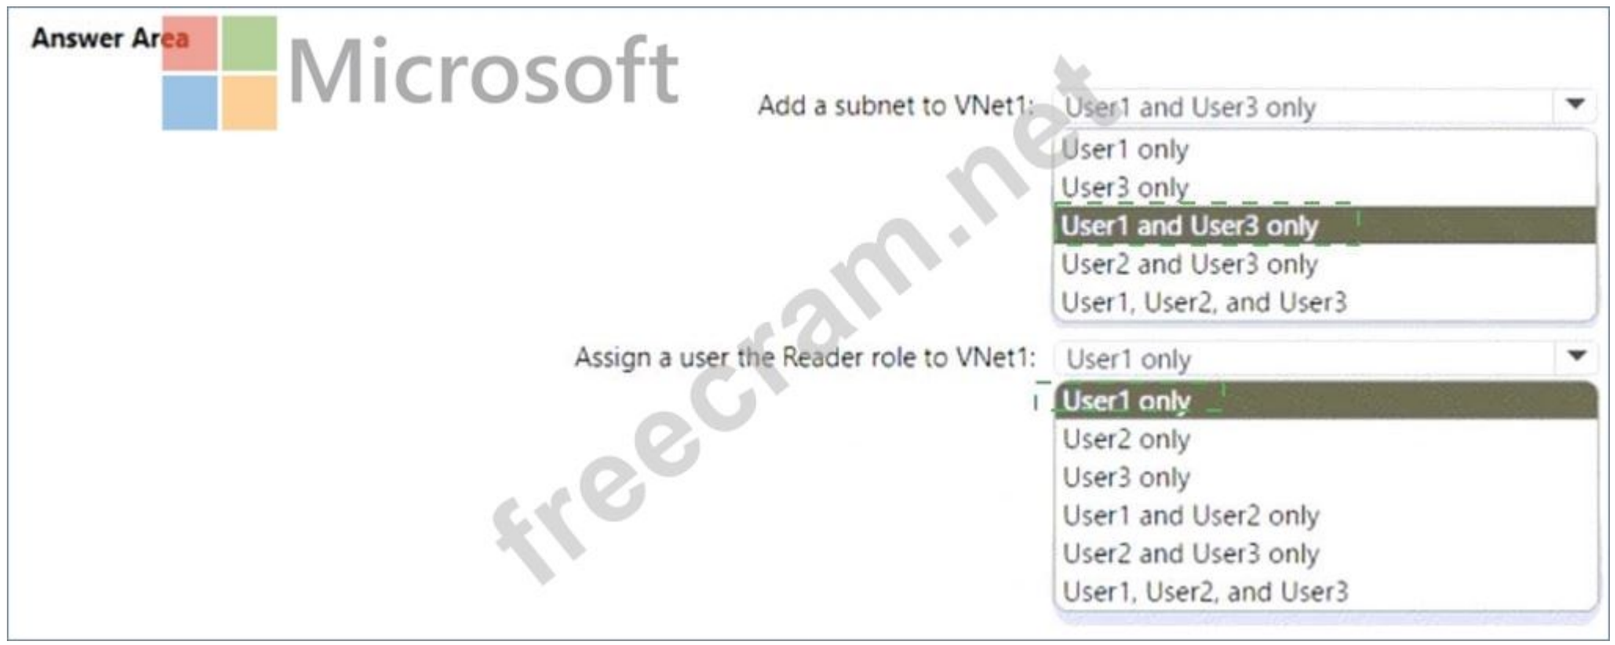

Explanation:

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

1. Understanding Each User's Role

User

Role

Key Capabilities

User1

Owner

Full access to all resources, including the ability to delegate access by assigning roles.

User2

Security Admin

Can manage security policies and settings (Azure Defender, Security Center) but cannot manage or configure networking resources.

User3

Network Contributor

Can manage networking resources, including VNets, subnets, and network interfaces, but cannot grant access or assign roles.

2. Task Analysis

Task 1: Add a subnet to VNet1

To add a subnet to a virtual network, the user must have permission to modify network resources.
The relevant Azure RBAC action is:
Microsoft.Network/virtualNetworks/subnets/write
Roles that include this permission:
Owner - has all permissions.
Network Contributor - can manage all network settings, including adding/removing subnets.
Therefore, User1 (Owner) and User3 (Network Contributor) can perform this task.
User2 (Security Admin) cannot, because that role does not include network configuration permissions.
The Answer: User1 and User3 only
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.
The Answer: 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
Correct Answer
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: 132
You have an Azure virtual network named VNet1 that contains a subnet named Subnet1. Subnet1 contains three Azure virtual machines. Each virtual machine has a public IP address.
The virtual machines host several applications that are accessible over port 443 to user on the Internet.
Your on-premises network has a site-to-site VPN connection to VNet1.

You discover that the virtual machines can be accessed by using the Remote Desktop Protocol (RDP) from the Internet and from the on-premises network.
You need to prevent RDP access to the virtual machines from the Internet, unless the RDP connection is established from the on-premises network. The solution must ensure that all the applications can still be accesses by the Internet users.
What should you do?
A. Modify the address space of the local network gateway.
B. Remove the public IP addresses from the virtual machines.
C. Modify the address space of Subnet1.
D. Create a deny rule in a network security group (NSG) that is linked to Subnet1.
Answer: (SHOW ANSWER)
You can filter network traffic to and from Azure resources in an Azure virtual network with a network security group. A network security group contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure resources.
You can use a site-to-site VPN to connect your on-premises network to an Azure virtual network. Users on your on-premises network connect by using the RDP or SSH protocol over the site-to-site VPN connection.
You don't have to allow direct RDP or SSH access over the internet. And this can be achieved by configuring a deny rule in a network security group (NSG) that is linked to Subnet1 for RDP / SSH protocol coming from internet.
Modify the address space of Subnet1 : Incorrect choice
Modifying the address space of Subnet1 will have no impact on RDP traffic flow to the virtual network.
Modify the address space of the local network gateway : Incorrect choice Modifying the address space of the local network gateway will have no impact on RDP traffic flow to the virtual network.
Remove the public IP addresses from the virtual machines : Incorrect choice If you remove the public IP addresses from the virtual machines, none of the applications be accessible publicly by the Internet users.
Reference:
https://docs.microsoft.com/en-us/azure/virtual-network/security-overview
https://docs.microsoft.com/en-us/azure/security/fundamentals/network-best-practices

NEW QUESTION: 133
You need to generate a shared access signature (SAS). The solution must meet the following requirements:
* Ensure that the SAS can only be used to enumerate and download blobs stored in container1.
* Use the principle of least privilege,
Which three settings should you enable? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

To generate a shared access signature (SAS) that meets the requirements, you should enable the following three settings:

Service: Blob

Allowed resource types: Container

Allowed permissions: Read and List

These settings will ensure that the SAS can only be used to enumerate and download blobs stored in container1, and not to perform any other operations on the storage account or the blobs. This follows the principle of least privilege, which means granting the minimum permissions necessary for a task.

You can use the Azure portal or Azure Storage Explorer to create a SAS token with these settings. For more information, see Create shared access signature (SAS) tokens for storage containers and blobs - Azure AI services | Microsoft Learn.

NEW QUESTION: 134

You have an Azure subscription that uses Azure Container Instances.

You have a computer that has Azure Command-Line Interface (CLI) and Docker installed.

You create a container image named image1.

You need to provision a new Azure container registry and add image1 to the registry.

Which command should you run for each requirement? To answer, select the options in the answer area NOTE: Each correct answer is worth one point.

Answer Area

Answer:

Explanation:

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

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

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

The documented workflow is:

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

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

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

Final Verified Answer:

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

NEW QUESTION: 135

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: 136

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: 137
You have a Recovery Service vault that you use to test backups. The test backups contain two protected virtual machines.
You need to delete the Recovery Services vault.
What should you do first?
A. From the Recovery Service vault, delete the backup data.
B. Modify the disaster recovery properties of each virtual machine.
C. Modify the locks of each virtual machine.
D. From the Recovery Service vault, stop the backup of each backup item.
Answer: (SHOW ANSWER)

NEW QUESTION: 138
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-N1')]"
  ],
  "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-N2')]"
        }
      ]
    }
  },
  {
    "apiVersion": "2017-03-30",
    "type": "Microsoft.Compute/virtualMachines",
    "name": "VM2",
    "zones": "2",
    "location": "EastUS2",
    "dependsOn": [
      "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-N1')]"
    ],
    "properties": {
      "hardwareProfile": {
        "vmSize": "Standard_A2_v2"
      },
      "osProfile": {
        "computerName": "VM2",
        "adminUsername": "AzureAdmin",
        "adminPassword": "[parameters('adminPassword')]"
      },
      "storageProfile": {
        "imageReference": "[variables('image')]",
        "osDisk": {
          "createOption": "FromImage"
        }
      },
      "networkProfile": {
        "networkInterfaces": [
          {
            "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-N2')]"
          }
        ]
      }
    }
  }
}
```

Answer Area

	Yes	No
VM1 and VM2 can connect to VNET1.	○	○
If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.	○	○
If the East US 2 region becomes unavailable, VM1 or VM2 will be available.	○	○

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: 139

You have an Azure subscription named Subscription1 that is used be several departments at your company.

Subscription1 contains the resources in the following table:

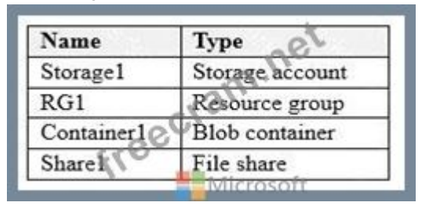

Name	Type
Storage1	Storage account
RG1	Resource group
Container1	Blob container
Share1	File share

Another administrator deploys a virtual machine named VM1 and an Azure Storage account named Storage2 by using a single Azure Resource Manager template.

You need to view the template used for the deployment.

From which blade can you view the template that was used for the deployment?

A. RG1

B. VM1

C. Storage1

D. Container1

Answer: (SHOW ANSWER)

1. View template from deployment history

Go to the resource group for your new resource group. Notice that the portal shows the result of the last deployment. Select this link.

2. You see a history of deployments for the group. In your case, the portal probably lists only one deployment. Select this deployment.

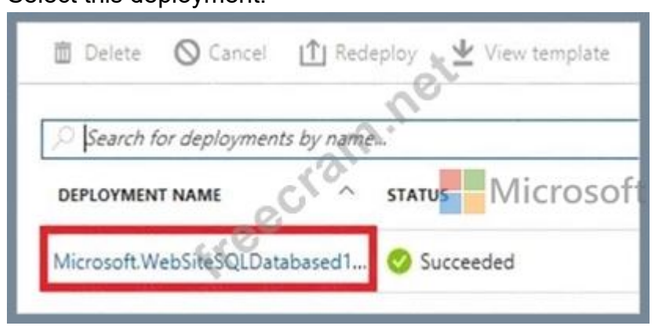

The portal displays a summary of the deployment. The summary includes the status of the deployment and its operations and the values that you provided for parameters. To see the template that you used for the deployment, select View template.

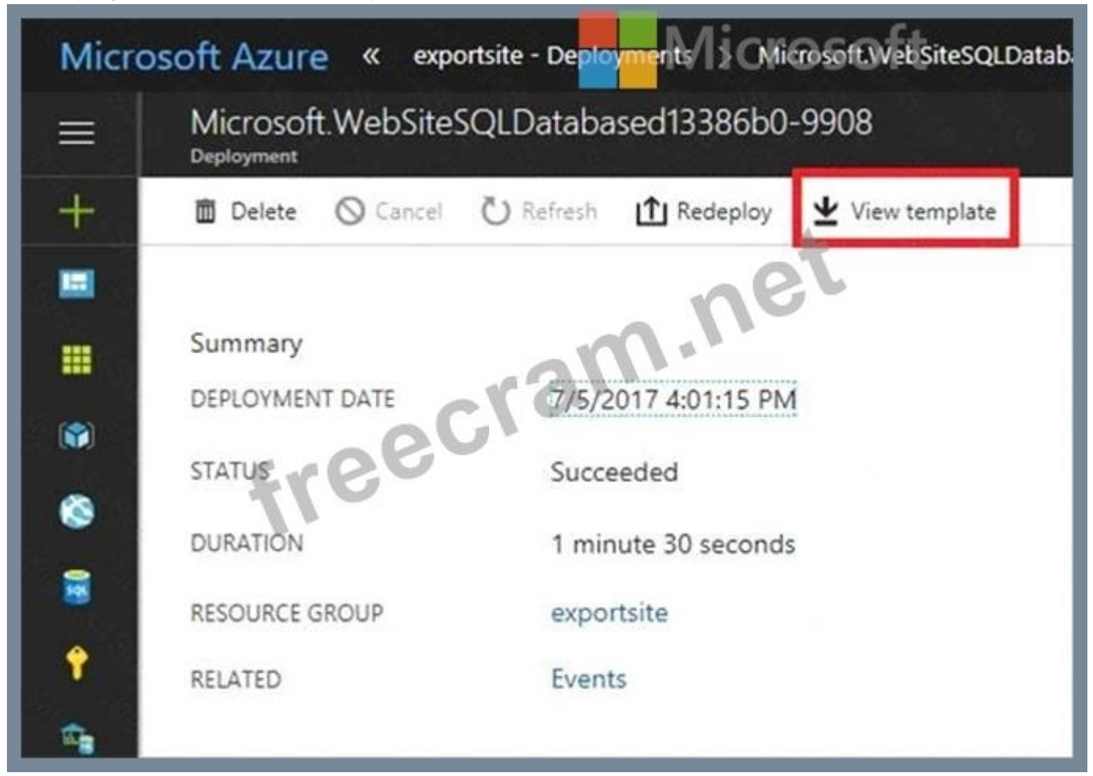

References: https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-manager-export- template

NEW QUESTION: 140

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:

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: 141

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: 142

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: 143
You deploy an Azure Kubernetes Service (AKS) cluster that has the network profile shown in the following exhibit.

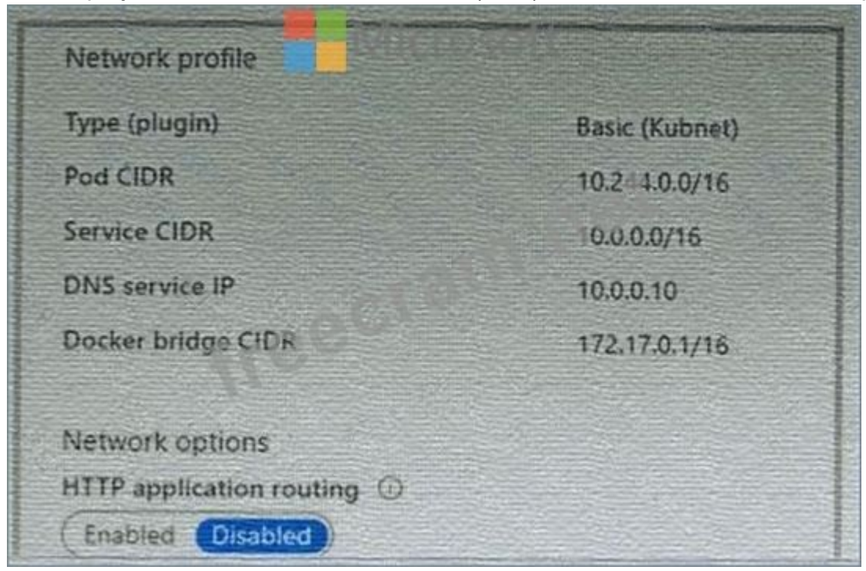

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

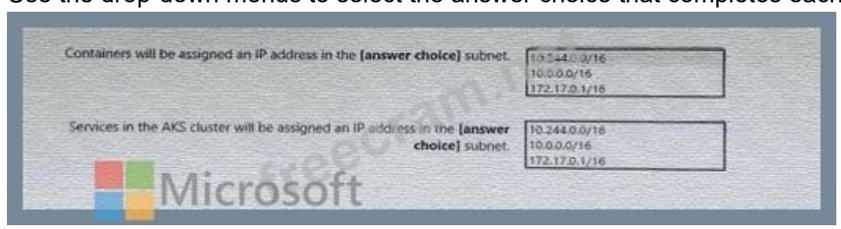

Answer:

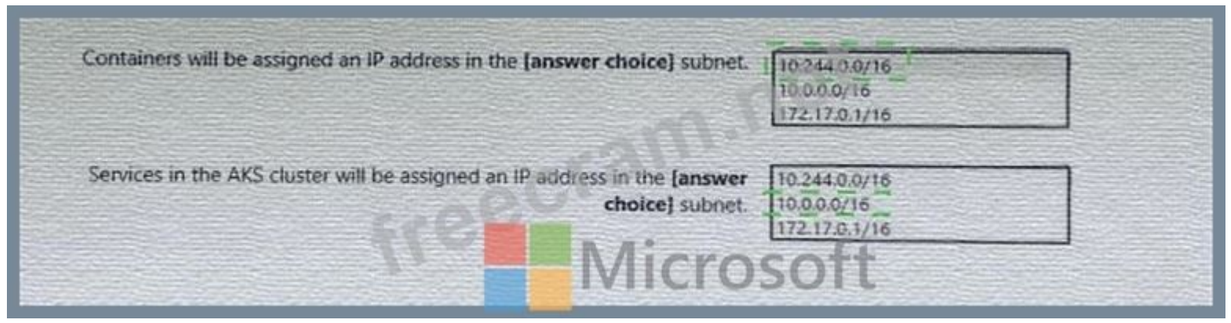

Explanation:

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: 144

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

You plan to deploy an Azure Monitor alert rule that will trigger an alert when CPU usage on VM1 exceeds 80 percent.

You need to ensure that the alert rule sends an email message to two users named User1 and User2.

What should you create for Azure Monitor?

A. an action group

B. a mail-enabled security group

C. a distribution group

D. a Microsoft 365 group

Answer: (SHOW ANSWER)

In Azure Monitor, an alert rule defines the condition (metric or log query) that triggers an alert, but notification and automation actions are managed through an Action Group.

An Action Group is a collection of notification preferences and actions used by Azure Monitor and Service Health alerts. It supports multiple actions such as:

* Email notifications

* SMS messages
* Push notifications
* Webhook triggers
* Automation runbooks
In this scenario, you must send an email message to two users (User1 and User2) whenever the CPU usage of VM1 exceeds 80%. To achieve this, you:
* Create an Action Group in Azure Monitor.
* Add two email receivers (User1 and User2) to that action group.
* Associate the Action Group with the alert rule monitoring VM1's CPU percentage.
Microsoft Official Documentation Extract (Paraphrased):
"Action groups are reusable notification collections that define how you want to be alerted when an alert rule triggers. You can specify multiple email recipients, SMS numbers, and webhook endpoints within one Action Group." (Source: Microsoft Learn - Create and manage action groups in Azure Monitor.)

NEW QUESTION: 145
You have an Azure subscription named Subscription1 that contains the following resource group:
Name: RG1
Region: West US
Tag: "tag1": "value1"
You assign an Azure policy named Policy1 to Subscription1 by using the following configurations:
Exclusions: None
Policy definition: Append tag and its default value
Assignment name: Policy1
Parameters:
- Tag name: Tag2
- Tag value: Value2
After Policy1 is assigned, you create a storage account that has the following configurations:
Name: storage1
Location: West US
Resource group: RG1
Tags: "tag3": "value3"
You need to identify which tags are assigned to each resource.
What should you identify? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:

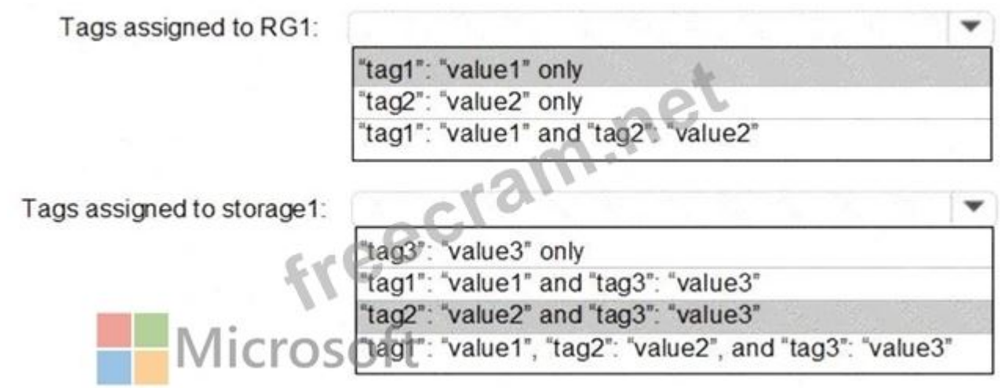

Box 1: " tag1 " : " value1 " only

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

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

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

NEW QUESTION: 146
You have an Azure Active Directory tenant named Contoso.com that includes following users:

Name	Role
User1	Cloud device administrator
User2	User administrator

Contoso.com includes following Windows 10 devices:

Name	Join type
Device1	Azure AD registered
Device2	Azure AD joined

You create following security groups in Contoso.com:

Name	Join type	Owner
Group1	Assigned	User1
Group2	Dynamic Device	User2

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
User1 can add Device2 to Group1	○	○
User2 can add Device1 to Group1	○	○
User2 can add Device2 to Group2	○	○

Answer:

Statements	Yes	No
User1 can add Device2 to Group1	○	○
User2 can add Device1 to Group1	○	○
User2 can add Device2 to Group2	○	○

Explanation:

Box 1: Yes
User1 is a Cloud Device Administrator.
Device2 is Azure AD joined.
Group1 has the assigned to join type. User1 is the owner of Group1.
Note: Assigned groups - Manually add users or devices into a static group.
Azure AD joined or hybrid Azure AD joined devices utilize an organizational account in Azure AD Box 2: No User2 is a User Administrator.
Device1 is Azure AD registered.
Group1 has the assigned join type, and the owner is User1.
Note: Azure AD registered devices utilize an account managed by the end user, this account is either a Microsoft account or another locally managed credential.
Box 3: Yes
User2 is a User Administrator.
Device2 is Azure AD joined.
Group2 has the Dynamic Device join type, and the owner is User2.
References:
https://docs.microsoft.com/en-us/azure/active-directory/devices/overview

NEW QUESTION: 147
You have an Azure subscription that contains the virtual networks shown in the following table.

Name	Location	Subnet
VNet1	East US	Subnet1, Subnet2
VNet2	West US	Subnet3

You add a service endpoint to each subnet as shown in the following table.

Subnet	Service endpoint
Subnet1	Microsoft.Storage.Global
Subnet2	Microsoft.KeyVault
Subnet3	Microsoft.Storage

You create the service endpoint policies shown in the following table.

Name	Location
Policy1	East US
Policy2	West US

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:

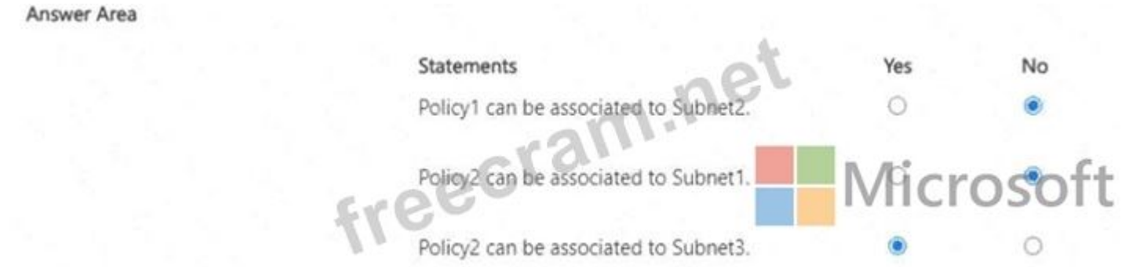

Service endpoint policies are used to restrict virtual network traffic over service endpoints to only specific Azure resources (for example, specific Storage accounts). Microsoft's service endpoint policy limitations state two key rules that apply directly here: (1) "Virtual networks must be in the same region and subscription as the service endpoint policy," and (2) "You can only apply a service endpoint policy on a subnet if service endpoints are configured for the Azure services listed in the policy." Microsoft Learn Given the configuration, Subnet2 has the service endpoint Microsoft.KeyVault, not Storage. Because the subnet does not have a Storage service endpoint configured, a Storage service endpoint policy (Policy1) can't be associated to Subnet2. Microsoft Learn Subnet1 is in VNet1 (East US), while Policy2 is created in West US. Since service endpoint policies must be in the same region as the virtual network/subnet they're applied to, Policy2 cannot be associated to Subnet1.
Microsoft Learn
Subnet3 is in VNet2 (West US) and has the Microsoft.Storage service endpoint configured, matching Policy2' s region and service requirement. Therefore, Policy2 can be associated to Subnet3. Microsoft Learn

NEW QUESTION: 148
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. Move VM1 to Subscription2.
B. Modify the IP address space of VNet2.
C. Provision virtual network gateways.
D. Move VNet1 to Subscription2.
Answer: (SHOW ANSWER)
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: 149
You have an Azure subscription that contains a resource group named RG1.
You plan to use an Azure Resource Manager (ARM) template named template1 to deploy resources. The solution must meet the following requirements:
* Deploy new resources to RG1.
* Remove all the existing resources from RG1 before deploying the new resources.
How should you complete the command? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

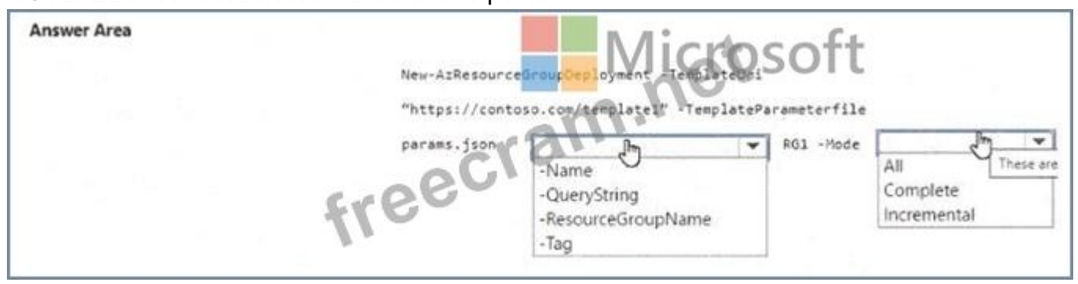

Answer:

Explanation:

Answer Area

https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-resourcegroupname
Specifies the name of the resource group to deploy.
https://learn.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroupdeployment?view=azps-9.3.0#-mode
Specifies the deployment mode. The acceptable values for this parameter are:
-Complete: In complete mode, Resource Manager deletes resources that exist in the resource group but are not specified in the template.
- Incremental: In incremental mode, Resource Manager leaves unchanged resources that exist in the resource group but are not specified in the template.

NEW QUESTION: 150
You have an Azure subscription.
You need to deploy a virtual machine by using an Azure Resource Manager (ARM) template.
How should you complete the template? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

```
{
    "$schema": "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
    ...
            "type": "Microsoft.Compute/virtualMachines",
            ...
            "dependsOn": [
                "[ [reference ▾] ('Microsoft.Network/networkInterfaces/', 'VM1')]"
            ],
            "properties": {
                "storageProfile": {
                    "[ [Array ▾] ":{
                        "publisher": "MicrosoftWindowsServer",
                        "Offer" : "WindowsServer",
                        "sku" : "2019-Datacenter",
                        "version" : "latest"
                    ...
}
```

reference
resourceId
Union

Array
Image
ImageReference
vhd

Answer:

Answer Area

```
{
    "$schema": "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#"
    ...
            "type": "Microsoft.Compute/virtualMachines",
            ...
            "dependsOn": [
                "[ [reference / resourceId / Union] ('Microsoft.Network/networkInterfaces/', 'VM1')]"
            ],
            "properties": {
                "storageProfile": {
                    "[ [Array / Image / ImageReference / vhd] ":{
                        "publisher": "MicrosoftWindowsServer",
                        "Offer" : "WindowsServer",
                        "sku" : "2019-Datacenter",
                        "version" : "latest"
                    ...
}
```

Explanation:

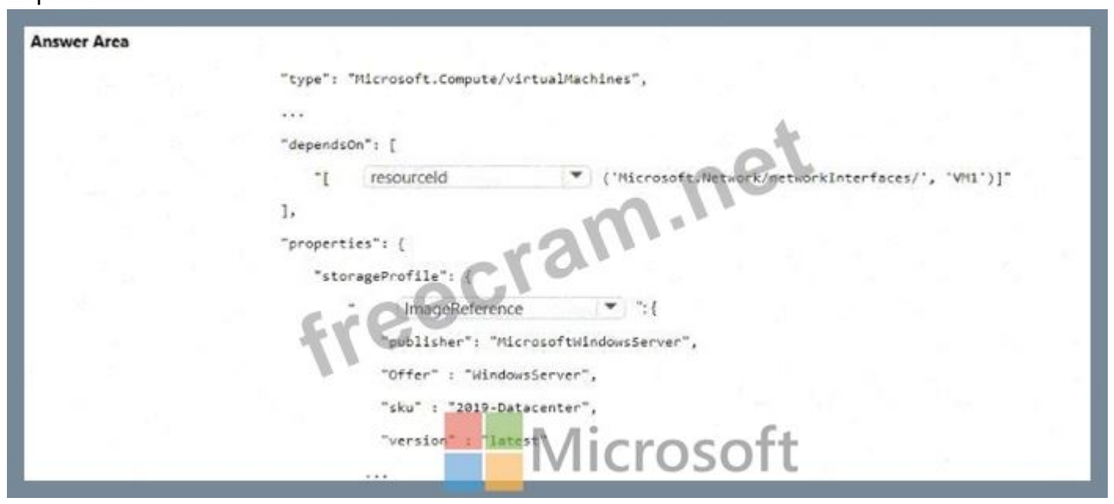

- dependsON: resoureceID

- storageProfile: ImageReference
Reference :
https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/resource-dependency#dependson
https://learn.microsoft.com/en-us/javascript/api/@azure/arm-compute/storageprofile?view=azure-node-latest

NEW QUESTION: 151
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: 152
You have the Azure virtual machines shown in the following table.

Name	IP address	Connected to
VM1	10.1.0.4	VNET1/Subnet1
VM2	10.1.10.4	VNET1/Subnet2
VM3	172.16.0.4	VNET2/SubnetA
VM4	10.2.0.8	VNET3/SubnetB

A DNS service is installed on VM1.

You configure the DNS server's settings for each virtual network as shown in the following exhibit

You need to ensure that all the virtual machines can resolve DNS names by using the DNS service on VM 1.

What should you do?

A. Configure peering between VNET1, VNET2, and VNET3.

B. Add service endpoints on VNET2 and VNET3.

C. Add service endpoints on VNET1.

D. Configure a conditional forwarder on VM1.

Answer: A (LEAVE A REPLY)

NEW QUESTION: 153

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: 154
You have an Azure subscription that contains a resource group named RG1.
You have a file named File1.bicep as shown in the File1.bicep exhibit. (Click the File1.bicep tab.)

```
resource vnet 'Microsoft.Network/virtualNetworks@2023-11-01' = {
  name: 'VNet1'
  location: resourceGroup().location
  tags: {
    CostCenter: '12345'
  }
  properties: {
    addressSpace: {
      addressPrefixes: [
        '10.0.0.0/24'
      ]
    }
    enableVmProtection: false
    enableDdosProtection: false
    subnets: [
      {
        name: 'Subnet1'
        properties: {
          addressPrefix: '10.0.0.0/24'
        }
      }
    ]
  }
}
```


You create a file named File2.bicep as shown in the File2.bicep exhibit. (Click the File2.bicep tab.)

```
resource vnet 'Microsoft.Network/virtualNetworks@2023-11-01' = {
  name: 'VNet1'
  location: resourceGroup().location
  tags: {
    CostCenter: '67890'
  }
  properties: {
    addressSpace: {
      addressPrefixes: [
        '10.0.0.0/16'
      ]
    }
    enableVmProtection: false
    enableDdosProtection: false
    subnets: [
      {
        name: 'Subnet2'
        properties: {
          addressPrefix: '10.0.1.0/24'
        }
      }
    ]
  }
}
```

You run the following PowerShell commands.

New-AzResourceGroupDeployment -ResourceOroupHame RGi -TeaplateFile Filel.bicep New-AzResourceGroupDeployment -Whatif -ResourceGroupNaae RG1 -TemplateFile File2.bicep 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
VNet1 has a `CostCenter` tag that has a value of 12345.	○	○
VNet1 has an IP address space of 10.0.0.0/16.	○	○
VNet1 has two subnets.	○	○

Answer:

Answer Area

Statements	Yes	No
VNet1 has a CostCenter tag that has a value of 12345.	○	○
VNet1 has an IP address space of 10.0.0.0/16.	○	○
VNet1 has two subnets.	○	○

Explanation:

Statements	Yes	No
VNet1 has a CostCenter tag that has a value of 12345.	●	○
VNet1 has an IP address space of 10.0.0.0/16.	○	●
VNet1 has two subnets.	○	●

NEW QUESTION: 155

You have a Microsoft Entra tenant named contoso.com.

You collaborate with an external partner named fabrikam.com.

You plan to invite users in fabrikam.com to the contoso.com tenant.

You need to ensure that invitations can be sent only to fabrikam.com users.

What should you do in the Microsoft Entra admin center?

A. From External collaboration settings, configure the Guest user access restrictions settings.

B. From Cross-tenant access settings, configure the Tenant restrictions settings.

C. From External collaboration settings, configure the Collaboration restrictions settings.

D. From Cross-tenant access settings, configure the Microsoft cloud settings.

Answer: (SHOW ANSWER)

Microsoft Entra ID provides External collaboration settings to control which external domains users can invite as guests. To restrict invitations so that only users from fabrikam.com can be invited, you must configure Collaboration restrictions.

Within the External collaboration settings, administrators can:

* Allow invitations only to users from specific domains
* Block invitations to all other external domains

Microsoft Entra documentation specifies that Collaboration restrictions are the control used to define allowed or blocked external domains for B2B guest invitations.

The other options do not meet the requirement:

* Guest user access restrictions control what guests can do after they are invited.
* Cross-tenant access - Tenant restrictions control inbound/outbound access behavior, not invitation eligibility.
* Microsoft cloud settings apply to cloud instances, not domain-based invitations.

Final Verified Answer:

C. From External collaboration settings, configure the Collaboration restrictions settings.

NEW QUESTION: 156
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 Redeploy blade, you click Redeploy.
Does this meet the goal?
A. Yes
B. No
Answer: (SHOW ANSWER)
When Azure schedules maintenance for a virtual machine, you can proactively move it to a new physical host by performing a self-service redeploy.
The Redeploy feature in the Azure portal allows you to:
* Move the VM to a new host node.
* Keep the same network interface, disks, and configuration.
* Resolve underlying host-level or platform issues proactively.
This action satisfies the requirement to move VM1 to a different host immediately and minimizes downtime.
This is explicitly documented in the Microsoft Learn - Redeploy Windows virtual machine to new Azure node guide.
Final Verified Answer: A. Yes

NEW QUESTION: 157
You have an Azure subscription.
Users access the resources in the subscription from either home or from customer sites. From home, users must establish a point-to-site VPN to access the Azure resources. The users on the customer sites access the Azure resources by using site-to-site VPNs.
You have a line-of-business app named App1 that runs on several Azure virtual machine. The virtual machines run Windows Server 2016.
You need to ensure that the connections to App1 are spread across all the virtual machines.
What are two possible Azure services that you can use? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.
A. a public load balancer
B. Traffic Manager
C. an Azure Content Delivery Network (CDN)
D. an internal load balancer
E. an Azure Application Gateway
Answer: (SHOW ANSWER)
Line of Business WebAPP works on VMs need internal load balancer. So D is needed. Then deploy WebAPP on VMs, check the link. https://docs.microsoft.com/en-us/azure/application-gateway/quick-create-portal So B is needed as well. The orignal answer is not accomplished.

NEW QUESTION: 158

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

NEW QUESTION: 159
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 Area

From the Azure portal:	Create an ExpressRoute circuit only. Create a virtual network gateway only. Create a virtual network gateway and a local network gateway. Create an ExpressRoute circuit and an on-premises data gateway. Create a virtual network gateway and an on-premises data gateway.
In the New York office:	Deploy ExpressRoute. Deploy a DirectAccess server. Implement a Web Application Proxy. Configure a site-to-site VPN connection.

Answer:

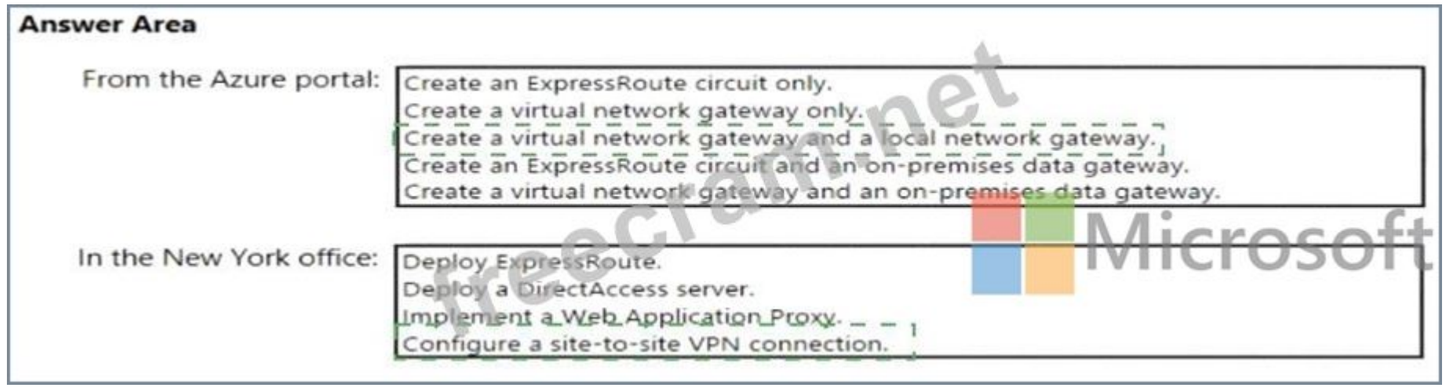

Explanation:

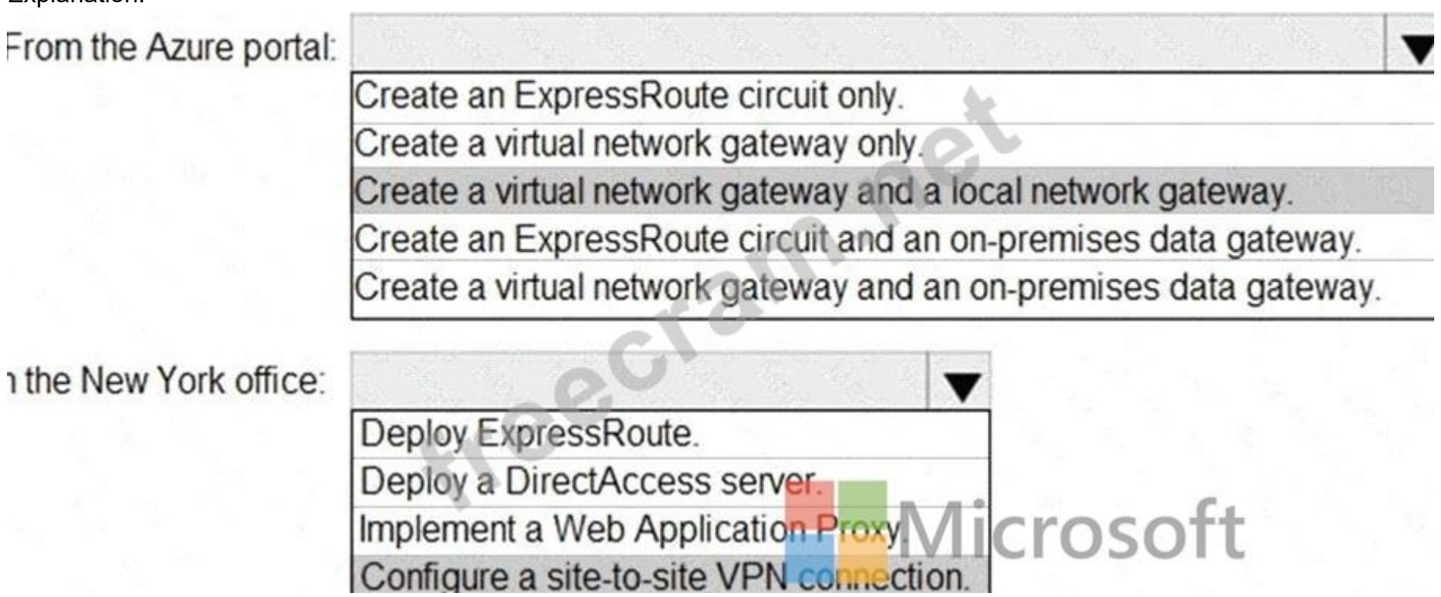

Box 1: Create a virtual network gateway and a local network gateway.

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.

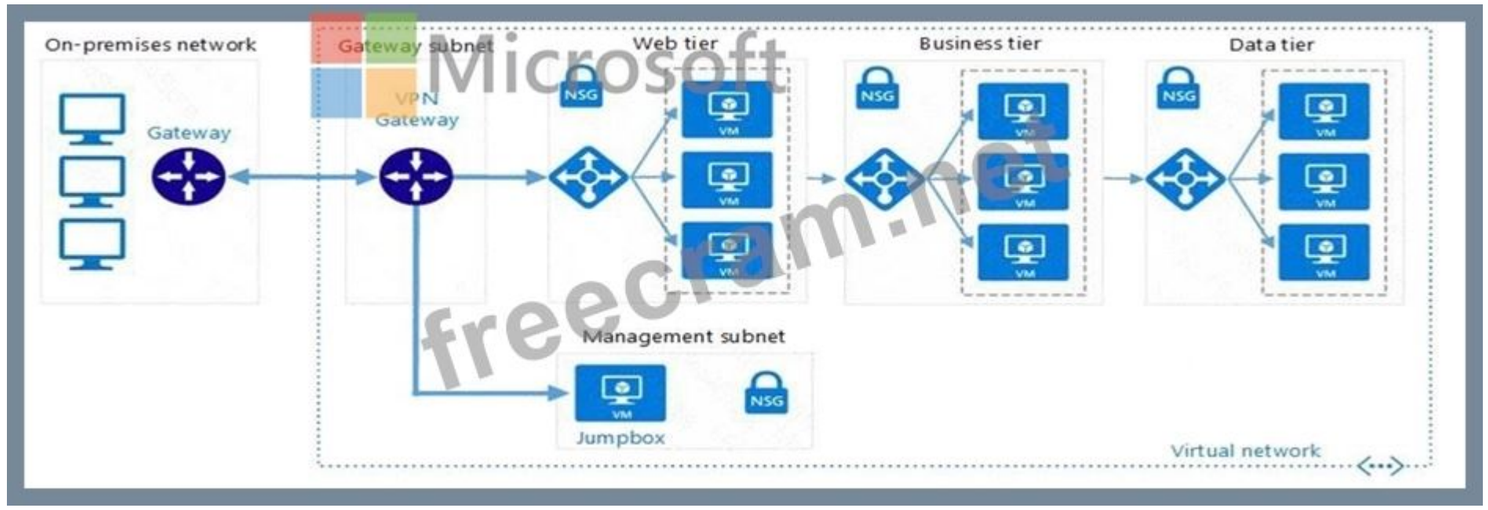

Scenario: Connect the New York office to VNet1 over the Internet by using an encrypted connection.

NEW QUESTION: 160

Your on-premises network contains an SMB share named Share1.
You have an Azure subscription that contains the following resources:
A web app named webapp1
A virtual network named VNET1
You need to ensure that webapp1 can connect to Share1.
What should you deploy?

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

NEW QUESTION: 161

You have an Azure web app named webapp1.
You have a virtual network named VNET1 and an Azure virtual machine named VM1 that hosts a MySQL database. VM1 connects to VNET1.
You need to ensure that webapp1 can access the data hosted on VM1.

What should you do?

A. Connect webapp1 to VNET1.

B. Deploy an internal load balancer.

C. Deploy an Azure Application Gateway,

D. Peer VNET1 to another virtual network.

Answer: (SHOW ANSWER)

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

In this scenario:

* VM1 hosts a MySQL database inside VNET1.

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

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

This ensures:

* Private, secure communication between the App Service and the virtual machine.

* No need to expose public IPs or deploy extra load balancers.

Incorrect options:

* B. Internal Load Balancer # Used to balance internal traffic among VMs, not for connecting web apps.

* C. Application Gateway # Provides HTTP load balancing, not private connectivity.

* D. VNet Peering # Used to connect different VNets, not to connect an App Service to a VNet.

Final Verified Answer: A. Connect webapp1 to VNET1

NEW QUESTION: 162

You have the App Service plans shown in the following table.

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

You plan to create the Azure web apps shown in the following table.

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

You need to identify which App Service plans can be used for the web apps.

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

NOTE: Each correct selection is worth one point.

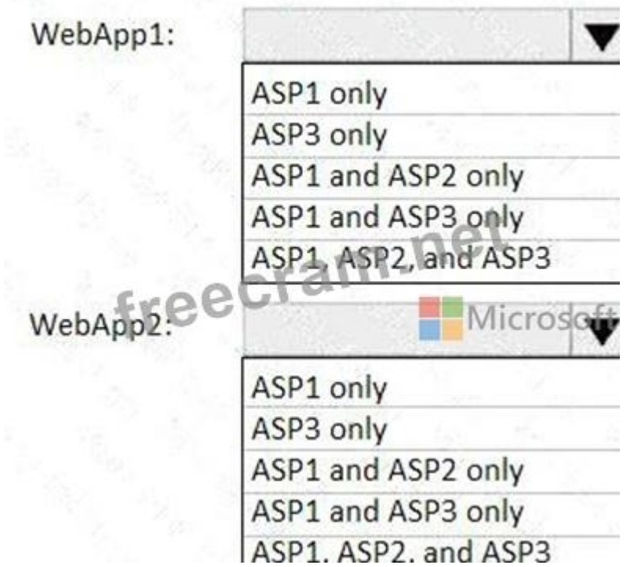

Answer:

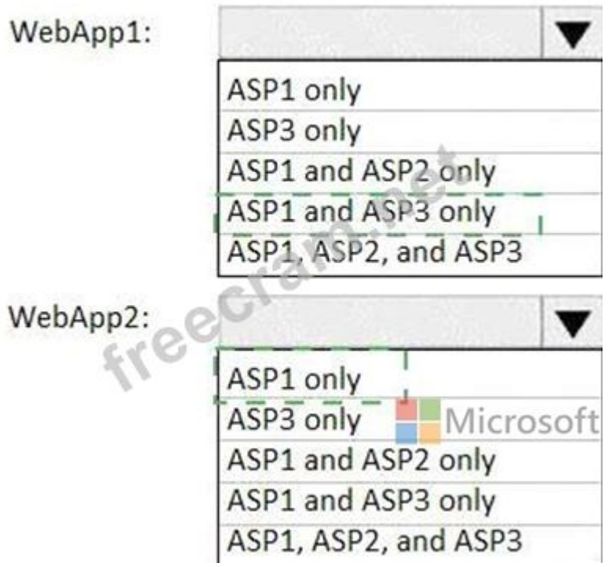

Explanation:

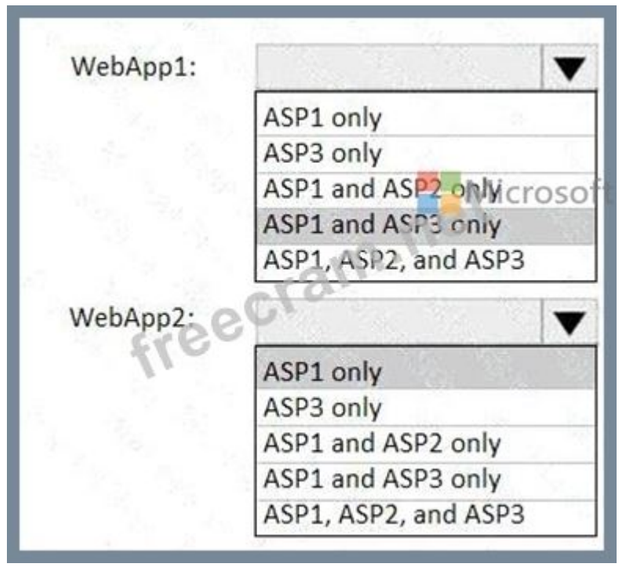

In Azure App Service, a Web App must be hosted in an App Service Plan (ASP) that meets three key compatibility conditions:
The App Service Plan and the Web App must be in the same region.
The operating system (Windows or Linux) of the App Service Plan must match the Web App type.
The runtime stack of the Web App must be supported on the App Service Plan's OS.
Let's apply these rules to the scenario:
App Service Plans:
Name
Operating System
Location
ASP1
Windows
West US
ASP2
Windows
Central US
ASP3
Linux
West US
Web Apps:
Name
Runtime Stack
Location
WebApp1
NET Core 3.0
West US
WebApp2

ASP.NET 4.7
West US
1## WebApp1 (.NET Core 3.0, West US)
Region requirement: Must use an App Service Plan in West US # # ASP1 and ASP3 qualify.
Runtime requirement: .NET Core 3.0 supports both Windows and Linux App Service Plans.
(Microsoft Learn: ".NET Core apps can run on both Windows and Linux App Service plans.")
Therefore, WebApp1 can use ASP1 (Windows, West US) or ASP3 (Linux, West US).
Answer: ASP1 and ASP3 only
2## WebApp2 (ASP.NET 4.7, West US)
Region requirement: Must be in West US # # ASP1 and ASP3 qualify.
Runtime requirement: ASP.NET 4.7 is a Windows-only framework; it cannot run on Linux App Service Plans.
(Microsoft Learn: "ASP.NET (non-Core) apps require a Windows-based App Service Plan.")
Therefore, only ASP1 (Windows, West US) is compatible.
Answer: ASP1 only
Final Verified Answers:
Web App
Compatible App Service Plans
WebApp1
ASP1 and ASP3 only
WebApp2
ASP1 only
Microsoft Documentation Extract (Azure App Service):
"App Service plans must be in the same region as the web app."
"Windows App Service plans host ASP.NET, .NET Core, and Node.js apps."
"Linux App Service plans host .NET Core, Node.js, Python, PHP, Java, and custom containers."
".NET Framework (ASP.NET 4.x) applications can only run on Windows-based App Service plans." Hence, the verified and Microsoft-official answer is:
WebApp1 # ASP1 and ASP3 only
WebApp2 # ASP1 only

NEW QUESTION: 163
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: 164
You configure the custom role 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:

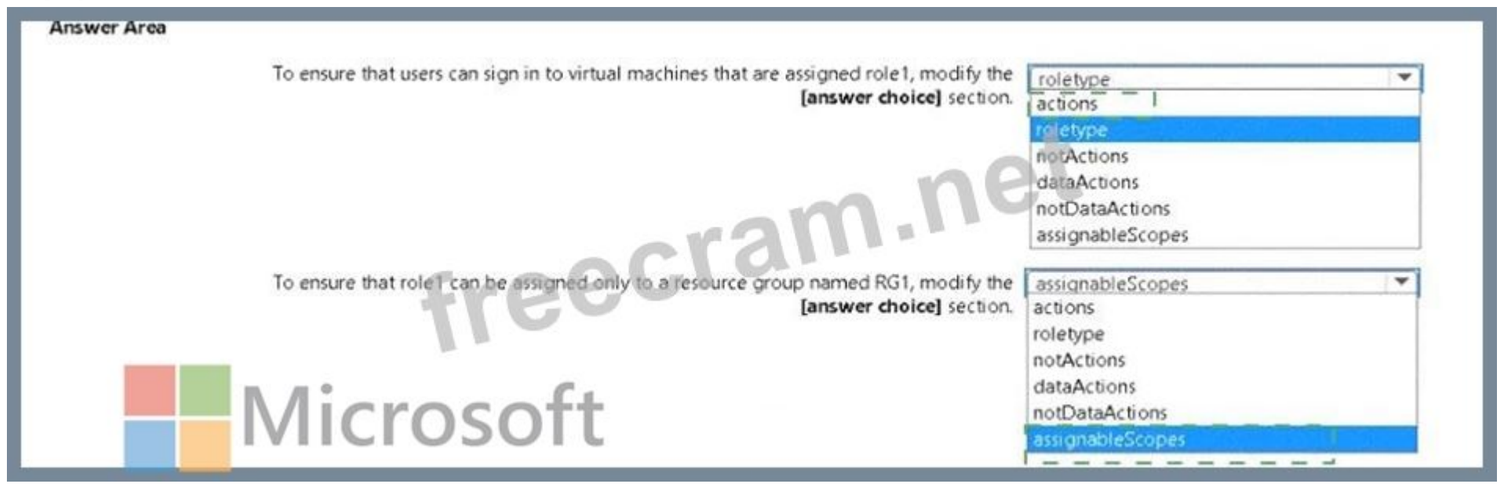

Explanation:

To ensure users can sign in to virtual machines that are assigned role1, modify the actions section.

To ensure that role1 can be assigned only to a resource group named RG1, modify the assignableScopes section.

In Azure Role-Based Access Control (RBAC), a custom role defines a collection of permissions that specify what actions a user or service principal can perform and the scope to which these permissions apply.

The JSON definition of a custom role contains key properties including:

* actions: Lists the management operations permitted (e.g., " Microsoft.Compute/virtualMachines/* " allows control over VMs).
* notActions: Explicitly denies certain operations that would otherwise be allowed.
* dataActions: Lists operations on data within a resource (e.g., blob reads/writes).
* assignableScopes: Defines the scopes (subscriptions, resource groups, or resources) where the custom role can be assigned.

According to Microsoft Azure Administrator documentation, if you want to ensure that users assigned a role can sign in to virtual machines, you must include the appropriate sign-in permissions in the actions section.

Specifically, you would add:

" Microsoft.Compute/virtualMachines/login/action "

This grants sign-in access to the VM via Azure RBAC.

Furthermore, if you want to restrict the scope of where the custom role can be assigned - for example, to a specific resource group like RG1 - you must define that restriction within the assignableScopes section.

For example:
" assignableScopes " : [
" /subscriptions/{subscriptionId}/resourceGroups/RG1 "
]
This ensures that the custom role is only visible and assignable within that resource group, preventing assignments at subscription or management group levels.
These definitions align with official Microsoft RBAC schema guidelines described in Azure Role Definitions and Custom Roles Overview for Azure Resource Manager.

NEW QUESTION: 165
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

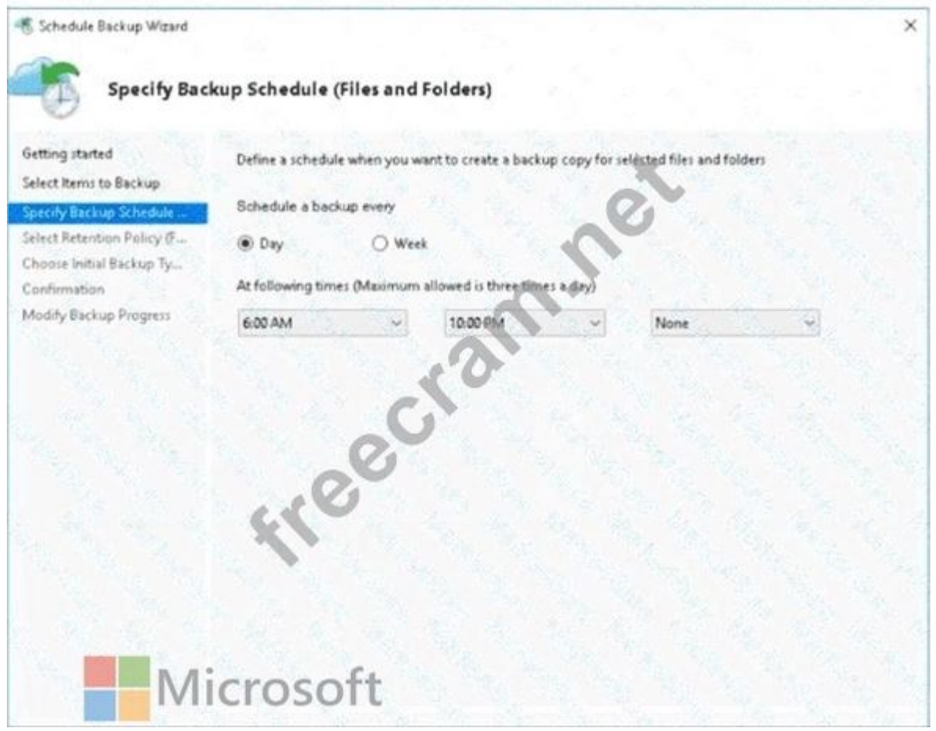

You plan to restore the backup to a different virtual machine.
You need to restore the backup to VM2.
What should you do first?
A. From VM1, install the Microsoft Azure Recovery Services Agent
B. From VM1, install the Windows Server Backup feature.
C. From VM2, install the Windows Server Backup feature.

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: 166

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.

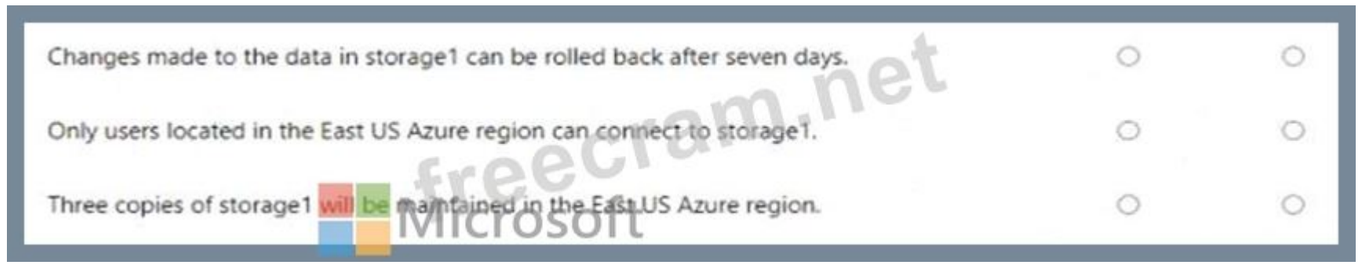

Answer:

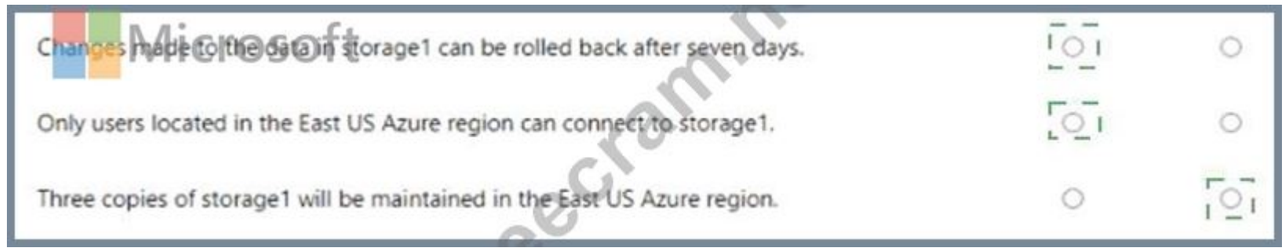

Explanation:

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: 167
You have an Azure Active Directory (Azure AD) tenant named contoso.onmicrosoft.com.
The User administrator role is assigned to a user named Admin1.
An external partner has a Microsoft account that uses the user1@outlook.com sign in.
Admin1 attempts to invite the external partner to sign in to the Azure AD tenant and receives the following error message: "Unable to invite user user1@outlook.com - Generic authorization exception." You need to ensure that Admin1 can invite the external partner to sign in to the Azure AD tenant.
What should you do?
A. From the Roles and administrators blade, assign the Security administrator role to Admin1.
B. From the Organizational relationships blade, add an identity provider.
C. From the Custom domain names blade, add a custom domain.
D. From the Users settings blade, modify the External collaboration settings.
Answer: (SHOW ANSWER)
In Azure AD, the ability to invite external (guest) users is controlled through External collaboration settings, found under Users # External collaboration settings in the Azure AD portal.
If a User administrator attempts to invite an external user but receives a "Generic authorization exception", this typically means the organization's external collaboration restrictions prevent user invitations.
According to Microsoft documentation (Configure external collaboration settings in Azure AD):
"To enable administrators or users to invite external guests, you must configure the External collaboration settings. You can define who can invite guest users into the directory, including administrators, users, or only specific roles." By updating these settings to allow User administrators to send invitations, Admin1 will be able to successfully invite the external partner (user1@outlook.com).

NEW QUESTION: 168
Yon have an Azure Storage account named storage1 that contains a blob container named comainer1. You need to prevent new content added to contalner1 from being modified for one year. What should you configure?
A. an access policy
B. the access level
C. the access tier
D. the Access control (JAM) settings
Answer: (SHOW ANSWER)
In Azure Storage, to ensure that newly added content cannot be modified or deleted for a fixed retention period, you must configure a container-level immutable storage policy, also known as a time-based retention policy or immutability policy. This feature is implemented through a stored access policy in Azure Blob Storage using the WORM (Write Once, Read Many) capability.
According to Microsoft's Azure Storage documentation, immutability policies can be configured in two ways:
Time-based retention policy - prevents data from being modified or deleted for a fixed duration (e.g., 1 year).
Legal hold - preserves data until the hold is explicitly cleared.
To implement the time-based retention policy, you configure a container-level access policy under Immutable blob storage settings. Once configured, any new blobs written to the container will be locked for the specified retention period and cannot be modified or deleted, ensuring compliance with retention regulations like SEC
17a-4(f) or GDPR.
The access level, access tier, or IAM settings do not provide this immutability. They control access permissions, performance, and user roles but do not enforce write-once protection.
Thus, the correct approach is to configure an access policy that enforces immutability for one year.

NEW QUESTION: 169
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 Area

Statements	Yes	No
A Site-to-Site connection can be established between VNET1 and VNET2.	○	○
VNET1 and VNET2 can be peered.	○	○
VNET1 and VNETA can be peered.	○	○

Answer:

Answer Area

Statements	Yes	No
A Site-to-Site connection can be established between VNET1 and VNET2.	○	○
VNET1 and VNET2 can be peered.	○	○
VNET1 and VNETA can be peered.	○	○

Explanation:

Answer Area

Statements	Yes	No
A Site-to-Site connection can be established between VNET1 and VNET2.	●	○
VNET1 and VNET2 can be peered.	●	○
VNET1 and VNETA can be peered.	○	●

In Microsoft Azure, connectivity between virtual networks depends on IP address space, region, and subscription constraints.
1## Site-to-Site VPN between VNET1 and VNET2 - YES:

Azure supports a Site-to-Site VPN connection between two VNets in different regions or subscriptions as long as both have non-overlapping address spaces and contain a VPN gateway. In this case, VNET1 (10.10.10.0/24 - West Europe) and VNET2 (172.16.0.0/16 - West US) have unique address spaces, so a S2S connection is possible.
2## VNET1 and VNET2 Peering - YES:
VNet peering enables direct communication over the Microsoft backbone network with low latency and no gateway requirement. Global VNet Peering supports VNets in different Azure regions (e.g., West Europe and West US) as long as address spaces do not overlap-which is true here.
3## VNET1 and VNETA Peering - NO:
VNETA's address space (10.10.128.0/17) overlaps with VNET1 (10.10.10.0/24), because both belong to the same 10.10.0.0/8 range. Azure explicitly blocks VNet peering between VNets with overlapping address spaces to avoid routing conflicts.
According to Microsoft Learn ("Virtual network peering - requirements and constraints"), peering is only supported when address spaces do not overlap. Therefore, VNET1 and VNETA cannot be peered.

NEW QUESTION: 170
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.

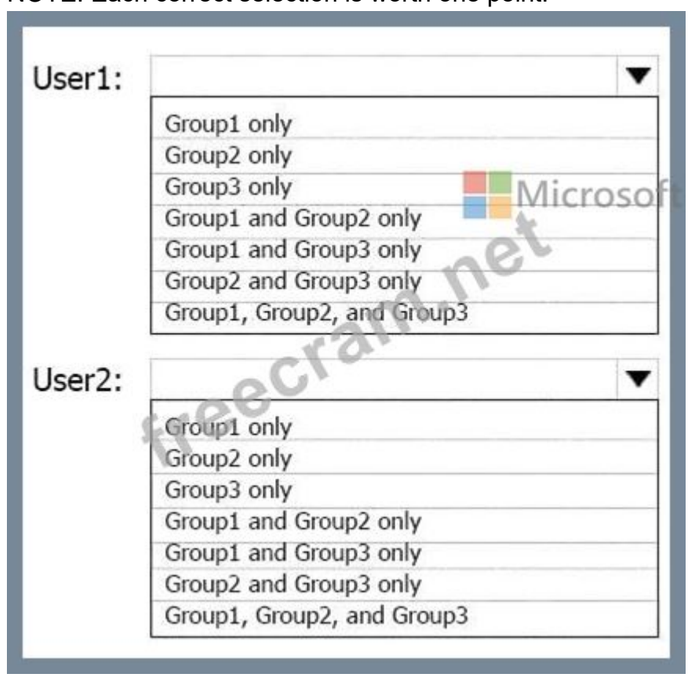

Answer:

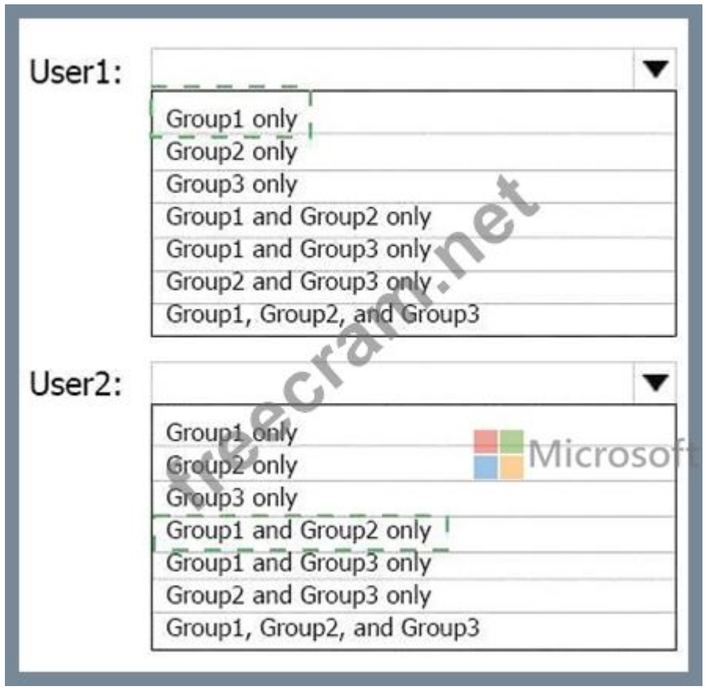

Explanation:

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: 171
You have an Azure subscription that contains a storage account named storage1.
You plan to create a blob container named contained.
You need to use customer-managed key encryption for contained.
Which key should you use?
A. an EC key that uses the P-384 curve only
B. an EC key that uses the P-521 curve only
C. an EC key that uses the P-384 curve or P-521 curve only
D. an RSA key with a key size of 4096 only
E. an RSA key type with a key size of 2048, 3072. or 4096 only

Answer: (SHOW ANSWER)

When configuring customer-managed keys (CMK) for Azure Storage encryption, Microsoft requires the use of Azure Key Vault-managed keys that meet specific cryptographic standards. According to the Azure Storage security and encryption documentation, only RSA keys are supported for customer-managed encryption keys. Elliptic Curve (EC) keys, regardless of curve type (P-384 or P-521), are not supported for Azure Storage CMK scenarios.

Azure Storage supports RSA keys with the following key sizes:

* 2048-bit
* 3072-bit
* 4096-bit

These key sizes align with Microsoft's encryption compliance and security baseline requirements. When a storage account is configured to use CMK, all supported services-including Blob containers-inherit encryption using the specified RSA key from Azure Key Vault.

Because the requirement is to encrypt the blob container contained using customer-managed keys, the only valid and supported choice is an RSA key with one of the supported key sizes listed above.

Final Answer: E. an RSA key type with a key size of 2048, 3072, or 4096 only

NEW QUESTION: 172

You have an Azure subscription that contains a storage account named storage 1.

You need to allow access to storage1 from selected networks and your home office. The solution must minimize administrative effort.

What should you do first for storage1?

A. Add a private endpoint.

B. Modify the Public network access settings.

C. Select Internet routing

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

Answer: (SHOW ANSWER)

To restrict access to a storage account in Azure to only specific networks or IP ranges (such as your home office), you must configure the firewall and virtual network settings under the networking section of the storage account.

Azure storage accounts, by default, are accessible from all networks. You can change this by modifying the Public network access setting to limit connections.

Here's the process per Microsoft Docs:

* Go to your Storage account # Networking # Firewalls and virtual networks.
* Under Public network access, select:
* "Enabled from selected virtual networks and IP addresses"
* Add:
* The IP address range of your home office.
* Any virtual networks (VNets) that should have access.

This allows you to control inbound traffic over Microsoft's backbone network without deploying additional infrastructure like Private Endpoints, which would increase administrative overhead.

Private endpoints (Option A) would indeed restrict access to a VNet via a private IP but require additional DNS and network configuration - hence not the minimal-effort option.

Access Control (IAM) (Option D) manages identity-based permissions (RBAC), not network-level access.

Therefore, modifying Public network access settings achieves the requirement with minimal effort.

Final Verified Answer: B. Modify the Public network access settings

NEW QUESTION: 173

You have an Azure subscription.

You have 100 Azure virtual machines.

You need to quickly identify underutilized virtual machines that can have their service tier changed to a less expensive offering.
Which blade should you use?
A. Metrics
B. Customer insights
C. Monitor
D. Advisor
Answer: (SHOW ANSWER)
The Advisor dashboard displays personalized recommendations for all your subscriptions. You can apply filters to display recommendations for specific subscriptions and resource types. The recommendations are divided into five categories:
Reliability (formerly called High Availability): To ensure and improve the continuity of your business-critical applications. For more information, see Advisor Reliability recommendations.
Security: To detect threats and vulnerabilities that might lead to security breaches. For more information, see Advisor Security recommendations.
Performance: To improve the speed of your applications. For more information, see Advisor Performance recommendations.
Cost: To optimize and reduce your overall Azure spending. For more information, see Advisor Cost recommendations.
Operational Excellence: To help you achieve process and workflow efficiency, resource manageability and deployment best practices. . For more information, see Advisor Operational Excellence recommendations.

NEW QUESTION: 174
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

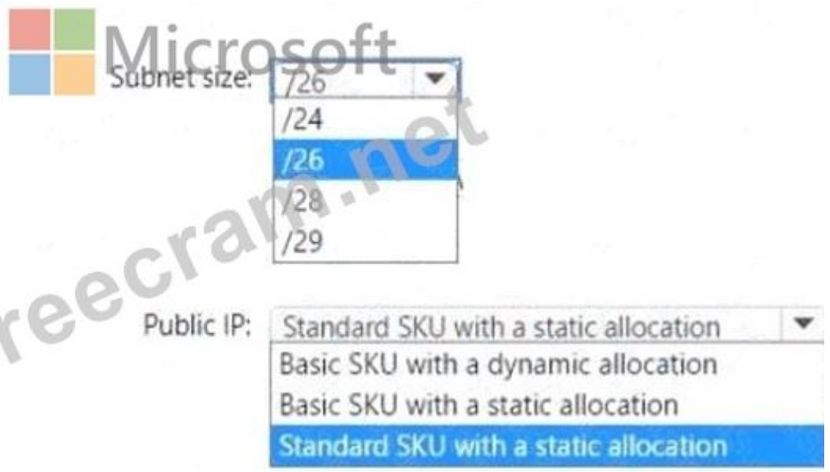

Answer:

Answer Area

Explanation:

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

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

1## Support for host scaling:

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

2## Support uploading and downloading files:

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

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

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

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

NEW QUESTION: 175
You have an Azure virtual machine named VM1 that runs Windows Server. The VM was deployed using default drive settings. You sign in to VM1 as a user named User1 and perform the following actions:
* Create files on drive C.
* Create files on drive D.
* Modify the screen saver timeout.
* Change the desktop background.
You plan to redeploy VM1. Which changes will be lost after you redeploy VM1?
A. the new files on drive D
B. the modified screen saver timeout
C. the new files on drive C
D. the new desktop background
Answer: (SHOW ANSWER)
When an Azure virtual machine is redeployed, Azure moves the VM to a new physical host. During this process, the OS disk and any attached data disks are preserved, but the temporary disk is not. On Windows VMs deployed with default settings, drive D: is the temporary disk, backed by local storage on the host.
Microsoft documentation states that the temporary disk is non-persistent and data on it can be lost during maintenance, redeploy, or deallocation events.
Files created on drive C: reside on the OS disk, which is persistent and retained across redeployments. User profile and system settings such as screen saver timeout and desktop background are stored on the OS disk and therefore remain intact after redeploy. Consequently, the only change that will be lost is data written to drive D:.

NEW QUESTION: 176
You have an Azure virtual network named VNet1 that connects to your on-premises network by using a site- to-site VPN. VMet1 contains one subnet named Subnet1.
Subnet1 is associated to a network security group (NSG) named NSG1. Subnet1 contains a basic internal load balancer named ILB1. ILB1 has three Azure virtual machines in the backend pool.
You need to collect data about the IP addresses that connects to ILB1. You must be able to run interactive queries from the Azure portal against the collected data.
What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

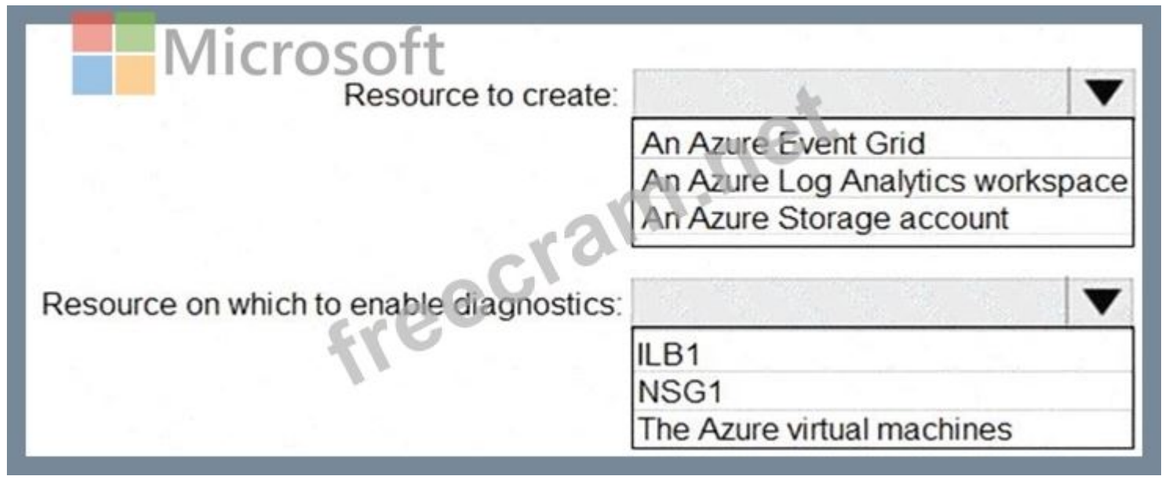

Answer:

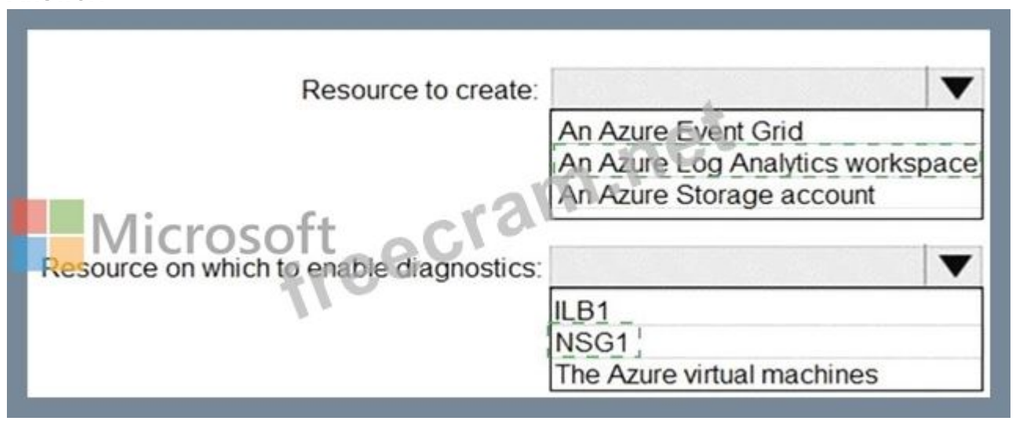

Explanation:

Box 1: An Azure Log Analytics workspace

In the Azure portal you can set up a Log Analytics workspace, which is a unique Log Analytics environment with its own data repository, data sources, and solutions.

Box 2: NSG1

NSG flow logs allow viewing information about ingress and egress IP traffic through a Network security group. Through this, the IP addresses that connect to the ILB can be monitored when the diagnostics are enabled on a Network Security Group.

We cannot enable diagnostics on an internal load balancer to check for the IP addresses.

As for Internal LB, it is basic one. Basic can only connect to storage account. Also, Basic LB has only activity logs, which doesn ' t include the connectivity workflow. So, we need to use NSG to meet the mentioned requirements.

NEW QUESTION: 177

You need to the appropriate sizes for the Azure virtual for Server2.

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

NOTE: Each correct selection is worth one point.

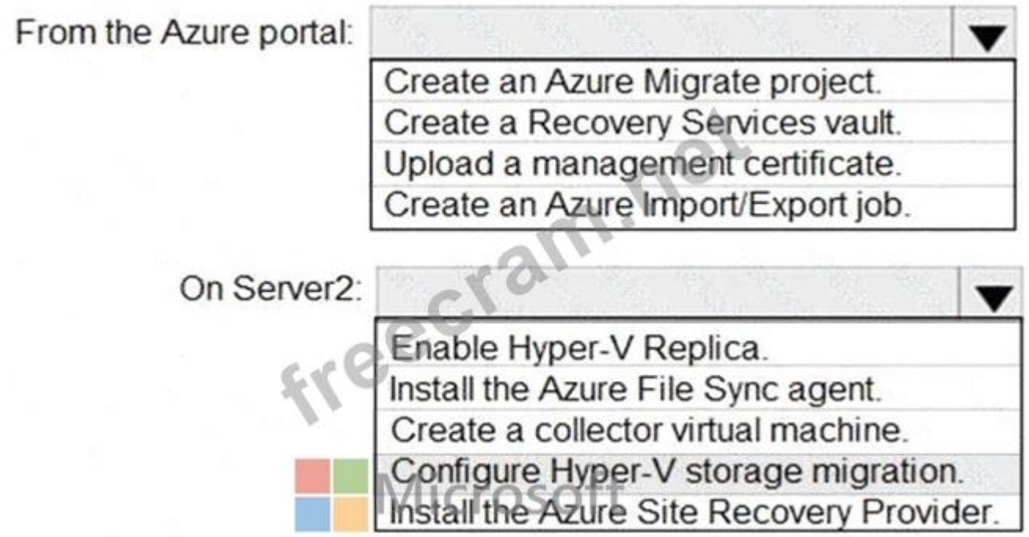

Answer:

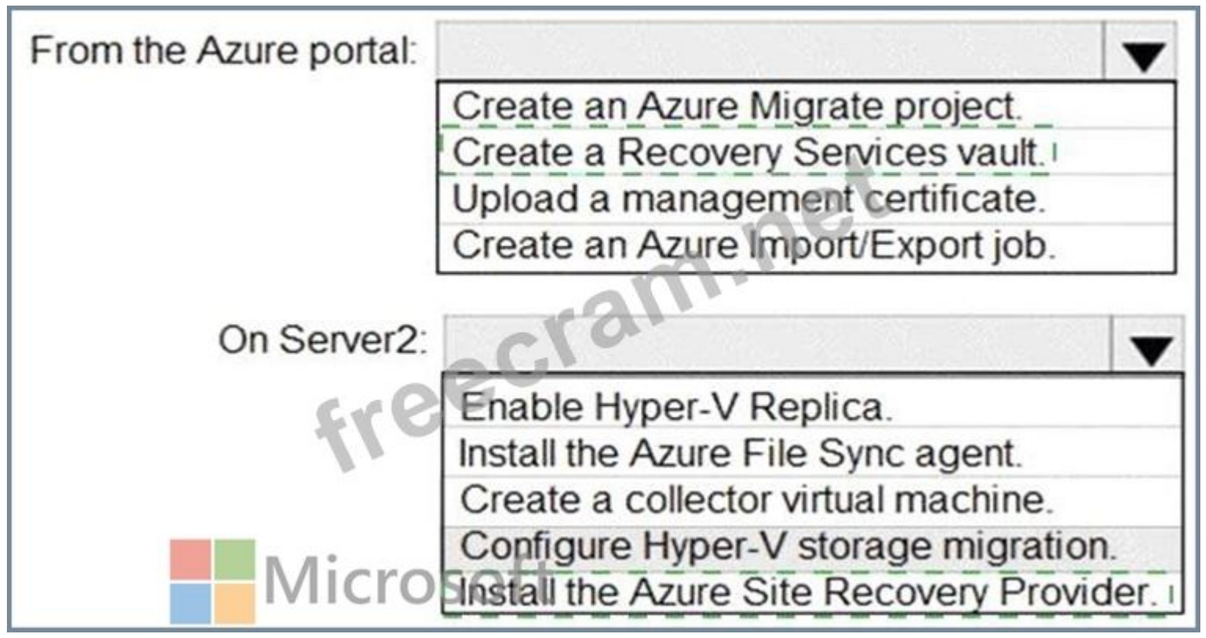

Explanation:

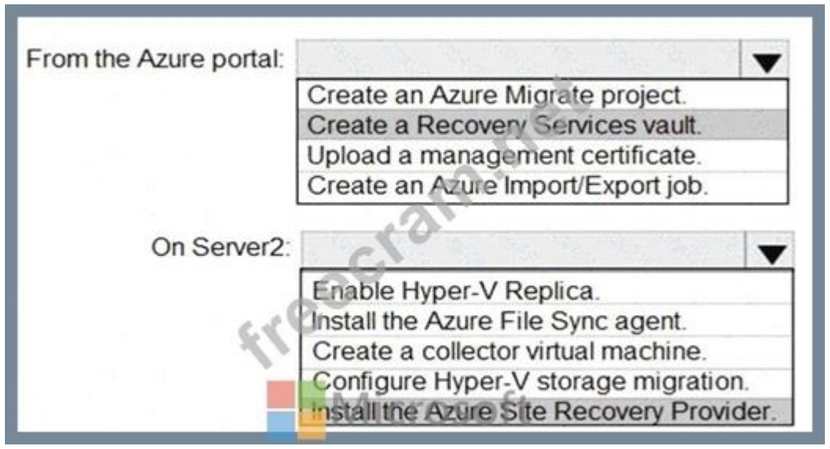

Box 1: Create a Recovery Services vault

Create a Recovery Services vault on the Azure Portal.

Box 2: Install the Azure Site Recovery Provider

Azure Site Recovery can be used to manage migration of on-premises machines to Azure.

Scenario: Migrate the virtual machines hosted on Server1 and Server2 to Azure.

Server2 has the Hyper-V host role.

References:

https://docs.microsoft.com/en-us/azure/site-recovery/migrate-tutorial-on-premises-azure

NEW QUESTION: 178

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.

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

Answer:

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

Explanation:

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

Statement 1: Yes

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

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

Statement 2: No

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

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

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

Reference:

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

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

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

NEW QUESTION: 179

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: 180

You have an Azure subscription that contains the virtual networks shown in the following table.
You need to ensure that all the traffic between VNet1 and VNet2 traverses the Microsoft backbone network.
What should you configure?

A. ExpressRoute
B. a private endpoint

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

NEW QUESTION: 181
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:

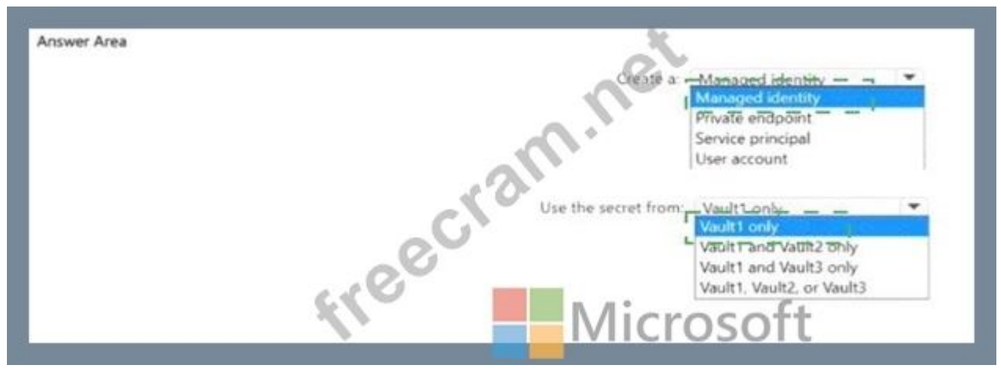

Explanation:

Valid AZ-104 Dumps shared by EduDump.com for Helping Passing AZ-104 Exam! EduDump.com now offer the **newest AZ-104 exam dumps**, the EduDump.com AZ-104 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AZ-104 dumps with Test Engine here: https://www.edudump.com/exams/Microsoft/AZ-104/premium/ (**454** Q&As Dumps, **35%OFF** **Special Discount Code: freecram**)