

Microsoft.AZ-104-KR.v2026-06-10.q199

□□□□:	AZ-104-KR
□□□□:	Microsoft Azure Administrator (AZ-104 Korean Version)
□□□:	Microsoft
□□ □□ □□□:	199
□□:	v2026-06-10
# □□ □:	131
# □□ □□□:	1990
https://www.krdump.com/Microsoft.AZ-104-KR.v2026-06-10.q199.html	

NEW QUESTION: 1

Azure □□□ □□□□.

□□□□ □□□ □□ □□□□□ □□□ □□□□ □□□□□□. □□□ □□□□ □□ □ □□□ □ VPN□ □□□□ Azure □□□□ □□□□□ □□□. □□ □□□□ □□□□ □□□ □ VPN□ □□□ □□□ Azure □□□□ □□□□□□.

□□ Azure □□ □□□□ □□□□ App1□□□□ LOB(□□ □□) □□ □□□□. □□ □□□ Windows Server 2016□ □□□□□.

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

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

□□: □□□ □□□ □□ 1□□ □□□ □□□□.

- A. □□ □□ □□□
- B. □□□ □□□
- C. Azure CDN(Content Delivery Network)
- D. □□ □□ □□□
- E. Azure □□□□□□ □□□□□

Answer: D,E (LEAVE A REPLY)

Line of Business WebAPP works on VMs need internal load balancer. So D is needed. Then deploy WebAPP on VMs, check the link. <https://docs.microsoft.com/en-us/azure/application-gateway/quick-create-portal> So B is needed as well. The original answer is not accomplished.

NEW QUESTION: 2

RG1□□□ □□□ □□□ □□□ Azure □□□ □□□□.

template1□□□ Azure Resource Manager(ARM) □□□□ □□□□ □□□□ □□□ □□□□□. □□□□ □□ □□ □□□ □□□□ □□□.

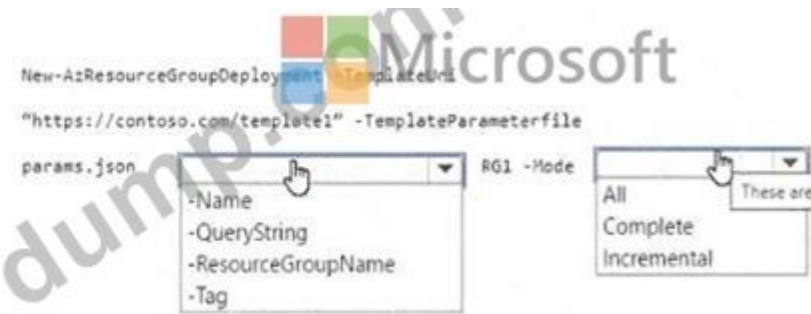
* RG1□ □□□ □□□□ □□□□□.

* □□□ □□□□ □□□□ □□ RG1□□ □□ □□□□ □□ □□□□□.

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

□□: □□ □□□ 1□□□□.

Answer Area



Answer:



Explanation:



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

Specifies the name of the resource group to deploy.

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

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

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

NEW QUESTION: 3

1. Azure Integration with Microsoft System Center Service Manager (ITSM) 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.

- A. ServiceNow.
- B. System Center Service Manager.
- C. ITSMC (ITSM) product or service.
- D. Provance.

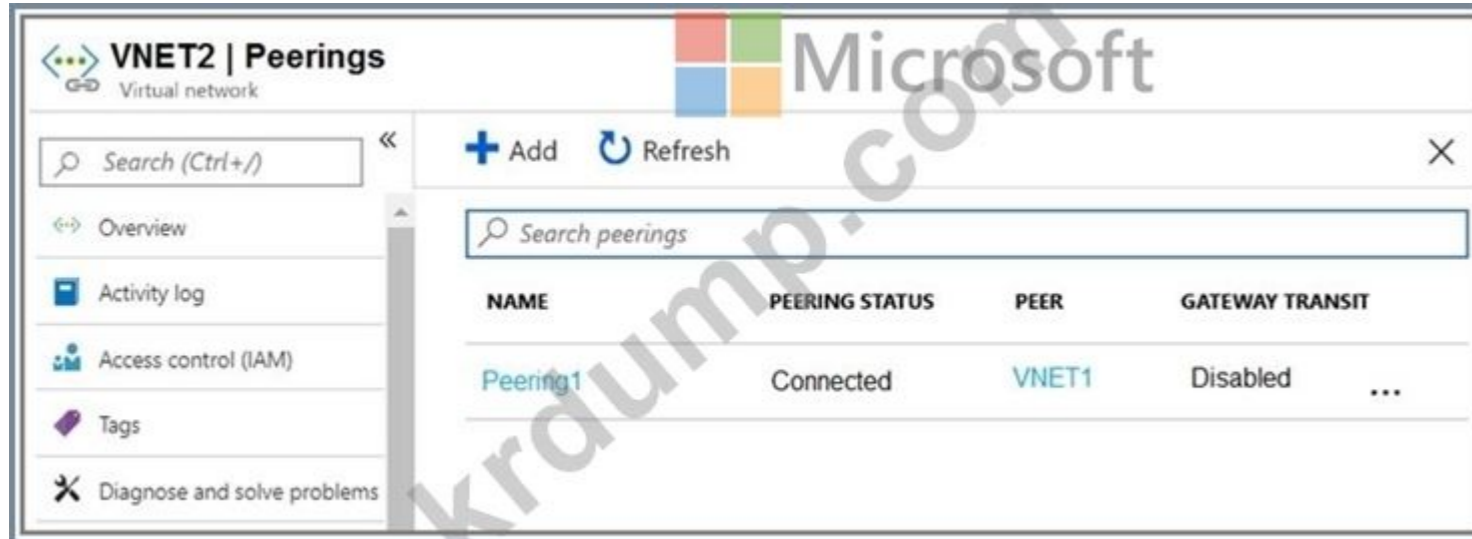
Answer: C (LEAVE A REPLY)

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

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

NEW QUESTION: 4

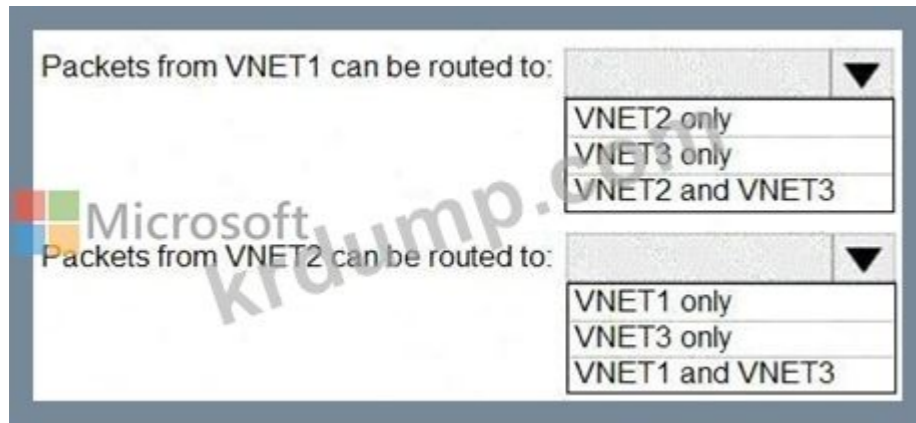
VNET2 is peered with VNET1.



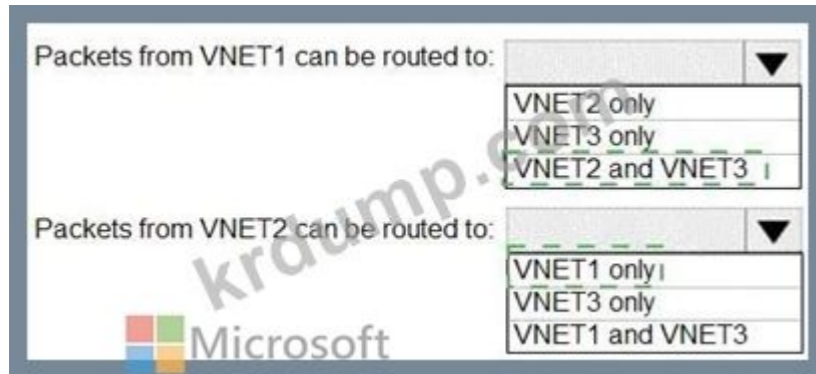
VNET3 is peered with VNET1.



What is the status of the peering between VNET2 and VNET3?
 The status is: Connected.



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: 5

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

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

Actions

- Configure the Diagnostic settings.
- Collect Windows performance counters from the Log Analytics agents.
- Create an alert rule.
- Create an Azure SQL database.
- Create a Log Analytics workspace.

Answer Area

The screenshot shows a question about disabling Gateway Transit and a list of five actions. The actions are: 1. Configure the Diagnostic settings. 2. Collect Windows performance counters from the Log Analytics agents. 3. Create an alert rule. 4. Create an Azure SQL database. 5. Create a Log Analytics workspace. The answer area is currently empty. Navigation arrows are visible on the right side of the actions list.

Answer:

The screenshot shows the same question and actions list, but with the answer area populated. The answer area contains three items: 1. Create a Log Analytics workspace. 2. Collect Windows performance counters from the Log Analytics agents. 3. Create an alert rule. A large watermark 'krdump.com' is overlaid on the image. Navigation arrows are visible on the right side of the answer area.

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: 6

Scenario: A company has a VNet1 in Azure. VNet1 contains two virtual machines, VM1 and VM2. The company wants to inspect all network traffic between VM1 and VM2 in VNet1 for a period of time.

Proposed solution: Configure Azure Network Watcher - Connection Monitor on VM1 and VM2. This solution meets the goal.

Proposed solution: Configure Azure Network Watcher - Packet Capture on VM1 and VM2. This solution meets the goal.

Proposed solution: Configure Azure Network Watcher - Connection Monitor on VM1 and VM2. This solution does not meet the goal.

Proposed solution: Configure Azure Network Watcher - Packet Capture on VM1 and VM2. This solution does not meet the goal.

Proposed solution: Configure Azure Network Watcher - Connection Monitor on VM1 and VM2. This solution does not meet the goal.

Proposed solution: Configure Azure Network Watcher - Packet Capture on VM1 and VM2. This solution does not meet the goal.

A.

B.

Answer: B (LEAVE A REPLY)

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: 7

Scenario: You are configuring Network Watcher on an Azure subscription. You need to ensure that the user Admin1 can enable Traffic Analytics. Admin1 is a member of the Azure AD(Azure Active Directory) group. You need to assign the minimum role to Admin1 to enable Traffic Analytics.

Options: A. Owner B. Contributor C. Reader D. Network Contributor

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: 8

Scenario: You are configuring Microsoft Entra ID for a new application. You need to ensure that the user Admin1 can manage the application. Admin1 is a member of the Azure AD(Azure Active Directory) group. You need to assign the minimum role to Admin1 to manage the application.

Options: A. Owner B. Contributor C. Reader D. Network Contributor

Answer: (SHOW ANSWER)

To manage the application, the user must have sufficient privileges to configure the application. As per Microsoft Azure documentation, the following built-in roles can manage the application:

Actions

Answer Area




Microsoft

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

Answer:

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



Microsoft

Explanation:

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

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

Step 1: Create a Group Naming Policy

Navigate to:

Microsoft Entra admin center # Groups # Naming policy.

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

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

Step 2: Set Add Prefix to Attribute

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

You can choose between two prefix types:

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

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

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

Step 3: Set Select Type to Department

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

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

This results in an automatically generated group name like:

FinanceSales, HRRecruiting, etc.

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

Official Microsoft Documentation Extract (Summarized):

"A naming policy can consist of prefixes or suffixes that include fixed strings or user attributes, such as [Department], [Company] , or [Office].

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

NEW QUESTION: 9

Azure .

Azure Storage .

Microsoft Azure Search resources, services, and docs (G+)

Home > Subscriptions > Subscription1 - Resources > New > Create storage account

Create storage account

✓ Validation passed

Basics Networking Advanced Tags **Review + create**

Basics

Subscription	Subscription1
Resource group	RG1
Location	(Europe) North Europe
Storage account name	storage16852
Deployment model	Resource manager
Account kind	StorageV2 (general purpose v2)
Replication	Locally-redundant storage (LRS)
Performance	Standard
Access tier (default)	Hot

Networking

Connectivity method	Private endpoint
Private Endpoint	(New) StorageEndpoint1 (blob) (privatelink.blob.core.windows.net)

Advanced

Secure transfer required	Enabled
Large file shares	Disabled
Blob soft delete	Disabled
Blob change feed	Disabled
Hierarchical namespace	Disabled
NFS v3	Disabled

Create < Previous Next >

[Download a template for automation](#)

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

□□: □□ □□□ 1□□□□.

Answer Area

The minimum number of copies of the storage account will be [answer choice].

3
1
2
3
4

Access tier (default)
Access tier (default)
Performance
Account kind
Replication

To reduce the cost of infrequently accessed data in the storage account, you must modify the [answer choice] setting.

Answer:

Answer Area

The minimum number of copies of the storage account will be [answer choice].

To reduce the cost of infrequently accessed data in the storage account, you must modify the [answer choice] setting.

Explanation:

Answer Area

The minimum number of copies of the storage account will be [answer choice].

To reduce the cost of infrequently accessed data in the storage account, you must modify the [answer choice] setting.

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: 10

VM1 is a virtual machine in an Azure subscription. You need to ensure that VM1 can be redeployed to a new physical host. You need to configure the VM1 ARM template to ensure that the VM1 can be redeployed to a new physical host. What should you do?

A. Set the `VM1 ARM1.json` file to `VM1 ARM1.json` in the Azure Resource Manager console.

B. Set the `VM1 ARM1.json` file to `VM1 ARM1.json` in the Azure Resource Manager console.

C. Set the `VM1 ARM1.json` file to `VM1 ARM1.json` in the Azure Resource Manager console.

D. Set the `VM1 ARM1.json` file to `VM1 ARM1.json` in the Azure Resource Manager console.

- A.
- B.

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: 11

VM1 is a virtual machine in an Azure subscription. You need to ensure that VM1 can be redeployed to a new physical host. You need to configure the VM1 ARM template to ensure that the VM1 can be redeployed to a new physical host. What should you do?

- A.
- B.
- C.
- D.

Answer: D (LEAVE A REPLY)

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: 12

VM1 is a virtual machine in an Azure subscription. You need to ensure that VM1 can be redeployed to a new physical host. You need to configure the VM1 ARM template to ensure that the VM1 can be redeployed to a new physical host. What should you do?

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

□□ □□ □□□ □□ □□□□ □□□ □□□□□□ □□□□□.

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

□□ □□ □□□ □□□ □□□□□ □□□ □□□□.

Name	Location
Policy1	East US
Policy2	West US

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

Answer Area

Statements	Yes	No
Policy1 can be associated to Subnet2.	<input type="radio"/>	<input type="radio"/>
Policy2 can be associated to Subnet1.	<input type="radio"/>	<input type="radio"/>
Policy2 can be associated to Subnet3.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
Policy1 can be associated to Subnet2.	<input type="radio"/>	<input checked="" type="radio"/>
Policy2 can be associated to Subnet1.	<input type="radio"/>	<input checked="" type="radio"/>
Policy2 can be associated to Subnet3.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
Policy1 can be associated to Subnet2.	<input type="radio"/>	<input checked="" type="radio"/>
Policy2 can be associated to Subnet1.	<input type="radio"/>	<input checked="" type="radio"/>
Policy2 can be associated to Subnet3.	<input checked="" type="radio"/>	<input type="radio"/>

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: 13

□□ □□□ □□ □□ □□□□.

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

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


link1
contoso.com

Save Discard Delete Access Control (IAM) Tags

Link name
link1

Link state
Completed

Provisioning state
Succeeded

Virtual network details
Virtual network id
/subscriptions/B3725433-236d-4361-b5ef-5b188fed87d0/resourceGroups/RG2/provi...

Virtual network
VNET2

Configuration
 Enable auto registration

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

Answer Area

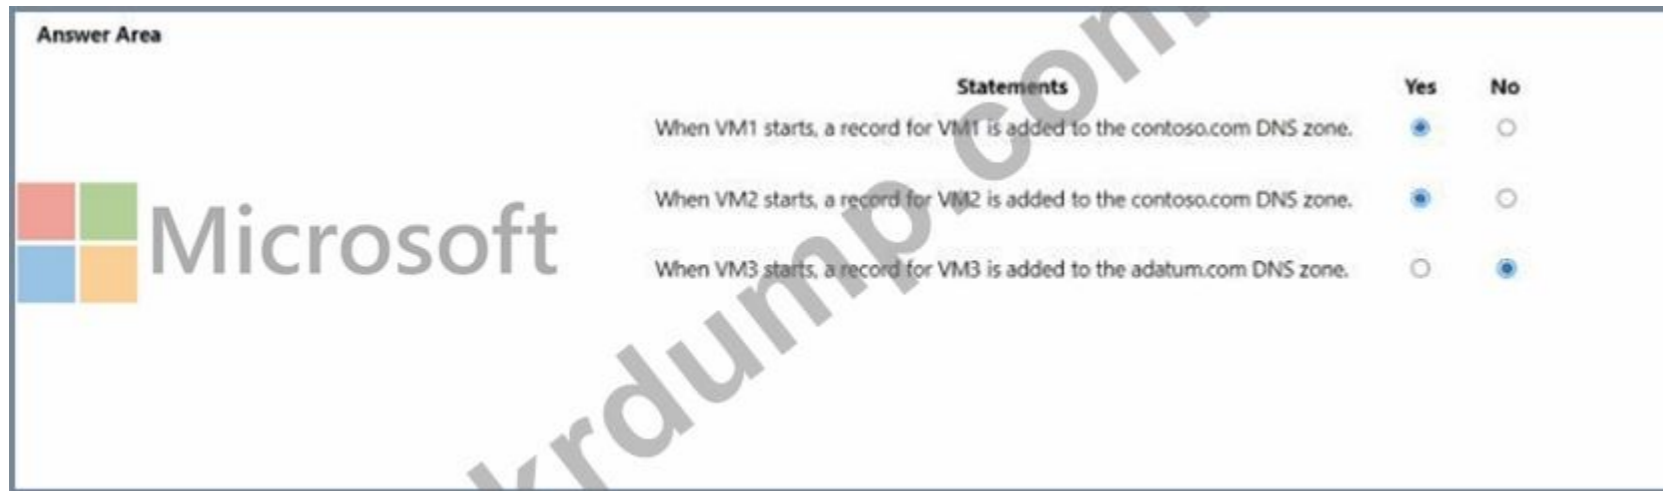
Statements	Yes	No
When VM1 starts, a record for VM1 is added to the contoso.com DNS zone.	<input type="radio"/>	<input type="radio"/>
When VM2 starts, a record for VM2 is added to the contoso.com DNS zone.	<input type="radio"/>	<input type="radio"/>
When VM3 starts, a record for VM3 is added to the adatum.com DNS zone.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
When VM1 starts, a record for VM1 is added to the contoso.com DNS zone.	<input checked="" type="radio"/>	<input type="radio"/>
When VM2 starts, a record for VM2 is added to the contoso.com DNS zone.	<input checked="" type="radio"/>	<input type="radio"/>
When VM3 starts, a record for VM3 is added to the adatum.com DNS zone.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:



All three VMs are in VNET2. Auto registration is enabled for private Azure DNS zone named contoso.com, which is linked to VNET2. So, VM1, VM2 and VM3 will auto-register their host records to contoso.com.

None of the VM will auto-register to the public Azure DNS zone named adatum.com. You cannot register private IPs on the internet (adatum.com) Box 1: Yes Auto registration is enabled for private Azure DNS zone named contoso.com.

Box 2: Yes

Auto registration is enabled for private Azure DNS zone named contoso.com.

Box 3: No

None of the VM will auto-register to the public Azure DNS zone named adatum.com Reference:

<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-name-resolution-for-vms-and-role-instances>

<https://docs.microsoft.com/en-us/azure/dns/private-dns-autoregistration>

<https://docs.microsoft.com/en-us/azure/dns/private-dns-virtual-network-links>

NEW QUESTION: 15

Scenario: You are configuring an Azure Storage account named storage1. You need to create a user named User1. You need to grant User1 the minimum permissions required to view the storage account properties and to read data within storage services such as blobs or files. You need to ensure that User1 can view the storage account properties and can read data within storage services such as blobs or files. You need to ensure that User1 can view the storage account properties and can read data within storage services such as blobs or files. You need to ensure that User1 can view the storage account properties and can read data within storage services such as blobs or files.

- A.
- B.

Answer: B (LEAVE A REPLY)

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

The Reader and Data Access role allows a user to:

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

* Microsoft.Storage/storageAccounts/regenerateKey/action

These permissions are included in roles such as:

* Storage Account Contributor

* Contributor

* Owner

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

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

NEW QUESTION: 16

VM1 and VM4 are located in different virtual networks (VNETs) in Azure. The requirement is to ensure that VM1 can communicate with VM4.

The solution must minimize administrative effort and cost.

A. VNET1 and VNET3 are connected via a virtual network gateway.

B. VM4 is assigned an IP address of 10.0.1.5/24.

C. VNET1 and VNET3 are connected via a virtual network peering.

D. A network security group (NSG) is created and associated with VM1 and VM4.

Answer: C (LEAVE A REPLY)

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

1. Background and Scenario Analysis

From the case study:

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

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

The solution must minimize administrative effort and cost.

2. Microsoft Documentation Insight: VNet Peering

According to Microsoft Learn: "Virtual network peering":

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

Enables private IP connectivity between resources across peered VNETs.

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

Provides low latency and high bandwidth.

Supports transitive routing through additional configurations.

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

3. Why the Other Options Are Incorrect

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

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

Without a connection path, a route has no effect.

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

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

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

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

4. Why Peering Is the Correct and Simplest Solution

"To deploy a container instance from an Azure container registry, the registry must be accessible either publicly with proper authentication (admin user or service principal with AcrPull permission) or privately using a Virtual Network with Private Link (Premium tier). If you receive authentication or access errors, the solution is to verify credentials or network accessibility, not to enable a dedicated data endpoint." In this case, the root cause of the deployment error is most likely related to image access authentication or tier limitations, not the use of dedicated data endpoints. Therefore, selecting Use dedicated data endpoint will not resolve the deployment failure.

The verified solution as per Microsoft Learn and AZ-104 exam content is to either:

Enable the admin user or

Assign a managed identity or service principal with the AcrPull role to the container instance.

Hence, the proposed solution does not meet the goal.

NEW QUESTION: 18

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

Name	Member of
User1	Group1
User2	Group2
User3	Group1

□□ □□ □□□ □□□ □□ □□□1□ □□□□□□ □□□ □□□ □□□□□ □□□□□.

Name	Type	Users to notify
Ingress	Metric	User1 and User3 only
Egress	Metric	User1 only
Delete storage account	Activity log	User1, User2, and User3
Restore blob ranges	Activity log	User1 and User3 only

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

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

□□: □□ □□□ 1□□□□.

Alert rules:

Microsoft

1
2
3
4

Action groups:

1
2
3
4

Answer:

Alert rules:

1
2
3
4

Action groups:

Microsoft

1
2
3
4

Explanation:

Box 1 : 4

As there are 4 distinct set of resource types (Ingress, Egress, Delete storage account, Restore blob ranges), so you need 4 alert rules. In one alert rule you can ' t specify different type of resources to monitor. So you need 4 alert rules.

Box 2 : 3

There are 3 distinct set of " Users to notify " as (User 1 and User 3), (User1 only), and (User1, User2, and User3). You can ' t set the action group based on existing group (Group1 and Group2) as there is no specific group for User1 only. So you need to create 3 action group.

Alert rules:

	▼
1	
2	
3	
4	

Action groups:

	▼
1	
2	
3	
4	

Reference:
<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/action-groups>

NEW QUESTION: 19

Scenario: A company has an Azure subscription. The company wants to create a Traffic Manager profile. The profile will be used to route traffic to a web application. The profile will be created in the East US region. The profile will have two endpoints. One endpoint will be in the East US region. The other endpoint will be in the West US region. The profile will be used to route traffic to the web application. The profile will be used to route traffic to the web application. The profile will be used to route traffic to the web application.

Question: Which role should be assigned to the user who will create the Traffic Manager profile?

Options: A. Traffic Manager Contributor B. Traffic Analytics Contributor C. Traffic Manager Contributor D. Admin1

Answer: B

A.

B.

Answer: B (LEAVE A REPLY)

The Traffic Manager Contributor role is not related to Traffic Analytics. Traffic Manager is a service that provides DNS-based load balancing and traffic routing across different regions and endpoints.

Traffic Manager Contributor is a role that allows you to create and manage Traffic Manager profiles, endpoints, and geographies1.

Traffic Analytics is a service that provides visibility into user and application activity in your cloud networks.

Traffic Analytics analyzes Azure Network Watcher network security group (NSG) flow logs to provide insights into traffic flow in your Azure cloud. With Traffic Analytics, you can visualize network activity, identify hot spots, secure your network, optimize your network deployment, and pinpoint network misconfigurations2.

To enable Traffic Analytics for an Azure subscription, you need to have a role that grants you the following permissions at the subscription level:

Microsoft.Network/applicationGateways/read

Microsoft.Network/connections/read
Microsoft.Network/loadBalancers/read
Microsoft.Network/localNetworkGateways/read
Microsoft.Network/networkInterfaces/read
Microsoft.Network/networkSecurityGroups/read
Microsoft.Network/publicIPAddresses/read
Microsoft.Network/routeTables/read
Microsoft.Network/virtualNetworkGateways/read
Microsoft.Network/virtualNetworks/read
Microsoft.OperationalInsights/workspaces/*

Some of the built-in roles that have these permissions are Owner, Contributor, or Network Contributor3.

However, these roles also grant other permissions that may not be necessary or desirable for enabling Traffic Analytics. Therefore, the best practice is to use the principle of least privilege and create a custom role that only has the required permissions for enabling Traffic Analytics4.

Therefore, to meet the goal of ensuring that an Azure AD user named Admin1 is assigned the required role to enable Traffic Analytics for an Azure subscription, you should create a custom role with the required permissions and assign it to Admin1 at the subscription level.

NEW QUESTION: 20

□□□ □□ □□□ □□□□ □□ □□□□ □□□□ □□□.

□□□ □□□□ □□□?

- A. □□□□□□ □□ □□(SSO) □ Active Directory □□□□□ □□□(AD FS)
- B. □□□□ □□ □□□ □ Single Sign-On(SSO)
- C. □□□□ □□ □□□ □□
- D. □□□□ □□ □ □□ □□□(SSO)

Answer: A (LEAVE A REPLY)

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: 21

□□: □ □□□ □□□ □□□□□ □□□□ □□□ □□ □ □□□□□. □ □□□ □□□ □□□ □□□ □ □□ □□□ □□□ □□□□ □□□□. □□ □□ □□□□ □□□ □ □□ □□ □□ □□, □□ □□ □□□□ □□□ □□ □ □□□□.

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

Adatum.com□□□ Microsoft Entra □□□□ Subscription1□□□ Azure □□□□ □□□□.

Adatum.com□□ Developers□□ □□□ □□□□. Subscription1□□ Dev□□ □□□ □□□ □□□□.

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

□□ □□: Dev□□ □□□ □□□□ Logic App Contributor □□□ □□□□□.

□□□ □□□ □□□□□?

- A. □
- B. □□□

Answer: A (LEAVE A REPLY)

In Microsoft Azure, Role-Based Access Control (RBAC) allows you to manage access to Azure resources by assigning roles to users, groups, and service principals. The Logic App Contributor role specifically grants the ability to create, edit, and manage Logic Apps but not assign permissions or modify access control (IAM).

According to the official Microsoft Azure documentation for built-in roles in Azure RBAC:

"The Logic App Contributor role lets you manage logic apps, but not access them. You can view, edit, update, and delete logic apps." In the given scenario, the Developers group needs the ability to create Azure Logic Apps within the Dev resource group. Assigning the Logic App Contributor role at the resource group level provides exactly that scope of control-users in the group can manage Logic App resources within that resource group without affecting other resources or permissions at the subscription level.

This satisfies the requirement under the principle of least privilege, as the Developers group gets only the permissions necessary to create and manage Logic Apps within Dev, and no more.

This is verified in the Microsoft Learn AZ-104 study guide under the topic "Manage Azure Role-Based Access Control (RBAC)", which emphasizes assigning the minimal level of role permissions required to complete administrative tasks at the appropriate scope (subscription, resource group, or resource).

Therefore, the proposed solution meets the goal because the Logic App Contributor role grants precisely the permissions needed.

NEW QUESTION: 22

Subscription1 is an Azure subscription.

Subscription1 has an alert rule named Alert1.

Alert1 is configured to send email notifications.

```
PS Azure:\> Get-AzureRmActionGroup

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

Alert1 is configured to send email notifications.

The email notifications are sent to the email address specified in the alert rule configuration.

The email address is: email@contoso.com.

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].

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



Explanation:

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



☐☐: ☐☐ ☐☐☐ 1☐☐☐☐.

Changes made to the data in storage1 can be rolled back after seven days.

Only users located in the East US Azure region can connect to storage1.

Three copies of storage1 will be maintained in the East US Azure region.

Answer:

Changes made to the data in storage1 can be rolled back after seven days.

Only users located in the East US Azure region can connect to storage1.

Three copies of storage1 will be maintained in the East US Azure region.

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: 24

Azure `ContReg1` `blob`.

`ContReg1` `blob` Azure `blob` `blob` `blob`.

`ContReg1` `blob` `blob` `blob` `blob` `blob` `blob` `blob` `blob`. `ContReg1` `blob` `blob` `blob` `blob`?

- A. `blob` `blob` `blob` `blob`.
- B. `blob` `blob`.
- C. `blob` `blob` `blob` `blob` `blob` `blob` `blob`.
- D. `blob` `blob` `blob` `blob`.

Answer: A (LEAVE A REPLY)

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: 25

`blob` `blob` `blob` Blob `blob` `blob` `blob` Sub1`blob` Azure `blob` `blob`.

Sub1`blob` User1`blob` User2`blob` `blob` `blob` `blob`. `blob` `blob` `blob` Sub1 `blob` `blob` `blob` `blob`.

`blob` `blob` `blob` `blob` Condition1`blob` `blob` `blob`.

```

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

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

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

```

□□ □□ □□□ □□ User1□ User2□ □□□ □□□□□.

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

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

ANSWER AREA

Statements	Yes	No
User1 can read blob2.	<input type="radio"/>	<input type="radio"/>
User1 can read blob3.	<input type="radio"/>	<input type="radio"/>
User2 can read blob1.	<input type="radio"/>	<input type="radio"/>

Answer:

ANSWER AREA

Statements	Yes	No
User1 can read blob2.	<input checked="" type="radio"/>	<input type="radio"/>
User1 can read blob3.	<input checked="" type="radio"/>	<input type="radio"/>
User2 can read blob1.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:


Answer Area

Statements

User1 can read blob2. Yes No

User1 can read blob3. Yes No

User2 can read blob1. Yes No



NEW QUESTION: 26

Azure .

Azure .

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

?

- A.
- B.
- C. Instance1 Instance2
- D. Instance3 Instance4

Answer: C (LEAVE A REPLY)

<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: 27

Azure .

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

Azure RG1 AF1 Azure .

AF1 ?

- A. VNET1
- B. VNET1 VNET2
- C. VNET1 VNET4
- D. VNET1, VNET2 VNET4
- E. VNET1, VNET2, VNET3 VNET4

Answer: (SHOW ANSWER)

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

From the table:

* VNET1 # West US # RG1 #

* VNET4 # West US # RG2 #

- * VNET2 # Central US #
- * VNET3 # Central US #

Resource group location does not restrict deployment-only region alignment matters.

Microsoft documentation confirms:

"Azure Firewall must be deployed into a virtual network in the same region."

NEW QUESTION: 28

☐☐ ☐☐☐ Azure ☐☐ ☐☐☐ ☐☐ ☐☐☐☐.

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

☐☐ ☐☐ ☐☐☐ ☐☐ ☐☐ ☐☐☐ Azure ☐☐☐ ☐☐☐☐☐.

Name	Management group
Subscription1	ManagementGroup21
Subscription2	ManagementGroup12

☐☐ ☐☐ ☐☐☐ Azure ☐☐☐ ☐☐☐☐.

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

☐☐ ☐ ☐☐☐ ☐☐, ☐☐☐ ☐☐☐ '☐'☐ ☐☐☐☐☐. ☐☐☐ ☐☐☐ '☐☐☐'☐ ☐☐☐☐☐. ☐☐: ☐☐☐ 1☐☐☐☐.

Answer Area

Statements	Yes	No
You can create a virtual network in Subscription1.	<input type="radio"/>	<input type="radio"/>
You can create a virtual machine in Subscription2.	<input type="radio"/>	<input type="radio"/>
You can move Subscription1 to ManagementGroup11.	<input type="radio"/>	<input type="radio"/>

Answer:
POWER TEST

Statements	Yes	No
You can create a virtual network in Subscription1.	<input type="radio"/>	<input checked="" type="radio"/>
You can create a virtual machine in Subscription2.	<input checked="" type="radio"/>	<input type="radio"/>
You can move Subscription1 to ManagementGroup11.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
You can create a virtual network in Subscription1.	<input type="radio"/>	<input checked="" type="radio"/>
You can create a virtual machine in Subscription2.	<input checked="" type="radio"/>	<input type="radio"/>
You can move Subscription1 to ManagementGroup11.	<input checked="" type="radio"/>	<input type="radio"/>

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: 29

☐☐ ☐☐☐ ☐☐ Microsoft Entra ☐☐☐☐ ☐☐☐☐☐☐☐☐.

* Azure Workspace! Azure Workspace? Azure Workspace.
* Windows Azure Workspace Azure Workspace.
* Azure Workspace Azure Workspace Azure Workspace.
Azure Workspace Azure Workspace Azure Workspace?

- A. Azure Monitor
- B. Windows Azure Workspace (WAD)
- C. Windows VM

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: 31

Azure: Azure Workspace Azure Workspace Azure Workspace. Azure Workspace Azure Workspace Azure Workspace. Azure Workspace Azure Workspace Azure Workspace.
Azure Workspace Azure Workspace Azure Workspace. Azure Workspace Azure Workspace Azure Workspace.
Azure Workspace Azure Workspace Azure Workspace Azure Workspace Azure Workspace. Azure Workspace Azure Workspace Azure Workspace Azure Workspace.
Azure Workspace Azure Workspace Azure Workspace Azure Workspace Azure Workspace Azure Workspace Azure Workspace Azure Workspace.
VM1 Windows Server VM1 VM2 Azure Workspace Azure Workspace Azure Workspace.
3 Azure VM1 VM2 Azure Workspace Azure Workspace Azure Workspace.
Azure: Azure Monitor Azure Workspace Azure Workspace Azure Workspace Azure Workspace Azure Workspace.
Azure Workspace Azure Workspace Azure Workspace?

- A.
- B.

Answer: B (LEAVE A REPLY)

Azure Monitor metrics such as Network In and Network Out only provide aggregated performance data - i.e., the total volume of bytes entering or leaving a network interface card (NIC). These metrics are useful for monitoring throughput, trends, or network utilization but do not capture or inspect packet-level data between two virtual machines.

To inspect network traffic between two Azure VMs (such as between VM1 and VM2), you must use Azure Network Watcher's Packet Capture or Connection Monitor feature.

According to the Microsoft Azure Administrator study guide and official documentation:

"Use Azure Network Watcher to monitor, diagnose, and view metrics for resources in a virtual network. The Packet Capture feature enables you to capture network traffic to and from a virtual machine and save it to an Azure Storage account for analysis." Therefore, simply creating metrics in Azure Monitor on Network In/Out will not meet the goal, because metrics do not provide detailed traffic inspection, only summary statistics.

To fulfill the requirement ("inspect all the network traffic from VM1 to VM2 for three hours"), the correct approach would be to use Network Watcher # Packet Capture, specifying VM1 as the target, capturing outbound traffic to VM2's IP, and setting the session duration for 3 hours.

Hence, the given solution does not meet the goal.

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: 33

Plan1 Azure App Service .

CPU 80% Plan1 . Plan1 ?

- A.
- B.
- C. () P1
- D. () S1
- E.

Answer: B (LEAVE A REPLY)

Azure App Service Plans determine the scaling behavior of web apps hosted on them. Scaling can be done manually or automatically depending on the pricing tier.

From the Microsoft Azure Administrator Study Guide and official documentation ("Scale instance count manually or automatically" - Microsoft Learn):

"To enable automatic scaling based on metrics such as CPU usage, memory percentage, or HTTP queue length, your App Service Plan must be in the Standard, Premium, PremiumV2, or higher tier.

You can then configure Scale Out (App Service plan) to use a Rules-Based method with performance thresholds." Available Scaling Options:

Manual: Fixed number of instances; no automatic scaling.

Automatic (Rules-Based): Create scaling rules based on metrics such as CPU > 80%, memory, or HTTP requests.

Scale Up (App Service Plan): Change pricing tier or hardware resources - does not provide auto-scaling.

In this scenario:

You already have a Standard plan (Plan1).

The requirement is to automatically scale out when CPU > 80%.

This behavior is achieved via Rules-Based scale-out, where you define a rule:

Metric: CPU Percentage

Condition: Greater than 80%

Action: Increase instance count

Therefore, you must choose Rules Based in the Scale out method settings for Plan1.

NEW QUESTION: 34

DCR1 ?

- A. WQL
- B. T-SQL
- C. XPath
- D. KQL

Answer: D (LEAVE A REPLY)

The planned change specifies that you must configure a Data Collection Rule (DCR) to collect only system events with Event ID 4648 from VM2 and VM4.

A Data Collection Rule (DCR) in Azure Monitor defines how data is collected from resources, filtered, and sent to destinations like Log Analytics workspaces. To define or query this data within Azure Monitor Logs or Log Analytics, you use Kusto Query Language (KQL).

From the Microsoft Learn: Azure Monitor Logs Documentation:

"Log queries in Azure Monitor are written in Kusto Query Language (KQL), the same query language used by Azure Data Explorer."

"KQL is optimized for querying large datasets, filtering by event IDs, sources, and event types." Other options:

- * WQL (WMI Query Language) - used for on-prem Windows event querying, not for Azure DCR.
- * T-SQL (Transact-SQL) - used for Azure SQL Database queries, not for monitoring data.
- * XPath - used in Event Viewer or XML-based event filtering, not within Azure Monitor DCR configuration.

Therefore, when you configure DCR1 to collect system events (Event ID 4648) from the specified VMs, the Kusto Query Language (KQL) is the correct and verified method to filter and process these events.

Example of a valid KQL expression for this requirement:

```
SecurityEvent
| where EventID == 4648
| where Computer in ("VM2", "VM4")
```

This aligns with the Azure Monitor and Log Analytics query methodology covered in AZ-104 official exam guide (Implement and manage monitoring).

NEW QUESTION: 35


Azure Storage Azure . vmimages . . ? . : 1 1 .

Answer Area




Answer:

Answer Area



Explanation:



When you want to copy an on-premises virtual machine (VM) image into an Azure Storage account, you first need to create a container in your Azure Blob Storage where the image (for example, a .vhd file) will reside.

According to Microsoft Azure Storage and AzCopy documentation, the AzCopy utility is a command-line tool used to copy data to and from Azure Storage accounts efficiently. It supports blob, file, and table services.

To create a new container in Azure Blob Storage using AzCopy, you use the make command. The make command creates a new container or directory in a Blob or File storage endpoint.

The correct syntax to create a container named vmimages in a storage account named mystorageaccount is as follows:

azcopy make 'https://mystorageaccount.blob.core.windows.net/vmimages' Explanation from Microsoft Documentation:

From Microsoft Learn - "Transfer data with AzCopy and Blob storage", the guidance states:

"Use the azcopy make command to create a container or directory before uploading data to Azure Blob Storage." azcopy make - creates the container (if it doesn't already exist).

https://mystorageaccount.blob.core.windows.net/vmimages - represents the destination URL where the container will be created.

Once created, you can then run another command such as:

azcopy copy 'C:\VMs\MyVM.vhd' 'https://mystorageaccount.blob.core.windows.net/vmimages' --blob-type PageBlob to upload the VHD file.

NEW QUESTION: 36

Sub1 is an Azure subscription, and RG1 is a resource group in Sub1. RG1 contains a Microsoft Entra ID group named HR. The HR group contains two users, User1 and User2. User1 is a member of the HR group, and User2 is not. The HR group is configured for self-service password reset (SSPR). What should you do to enable SSPR for only User1 while keeping it disabled for User2?

A. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR. Remove User2 from the HR group.

B. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR. Add User2 to the new group.

C. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

D. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR. Add User2 to the new group.

E. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

F. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR. Add User2 to the new group.

G. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

H. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

I. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

J. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

K. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

L. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

M. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

N. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

O. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

P. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

Q. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

R. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

S. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

T. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

U. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

V. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

W. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

X. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

Y. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

Z. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

AA. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

AB. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

AC. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

AD. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

AE. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

AF. Create a new Microsoft Entra group and add User1 to the group. Configure the group for SSPR.

Answer: C (LEAVE A REPLY)

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

To enable SSPR for only User1 while keeping it disabled for User2, the correct first action is to create a Microsoft Entra group and place User1 in that group. Microsoft Entra ID supports enabling SSPR for None, Selected, or All users. When Selected is used, SSPR is enabled by choosing a Microsoft Entra group; Microsoft's configuration flow explicitly states to select Selected, then browse for and select the Microsoft Entra group that contains the users enabled for SSPR.

Therefore, a group is required before targeted SSPR enablement can be applied. Authentication methods policies define which methods users can use, such as phone, email, or authenticator app, but they do not by themselves scope SSPR to only User1. Authentication context is used with Conditional Access, not SSPR targeting. Security defaults are tenant-wide baseline protections and are not the first step for enabling SSPR for one selected user. This maps directly to the AZ-104 study guide objective Manage Microsoft Entra users and groups, specifically Create users and groups and Configure self-service password reset (SSPR).

NEW QUESTION: 37

Sub1 is an Azure subscription, and RG1 is a resource group in Sub1. RG1 contains a Microsoft Entra ID group named HR. The HR group contains two users, User1 and User2. User1 is a member of the HR group, and User2 is not. The HR group is configured for self-service password reset (SSPR). What should you do to enable SSPR for only User1 while keeping it disabled for User2?

A. Sub1 Departments RG1 Department =HR

B. RG1 Department =HR

C. Azure Departments RG1

D. Departments RG1

Answer: C (LEAVE A REPLY)

NEW QUESTION: 38

Sub1 is an Azure subscription, and RG1 is a resource group in Sub1. RG1 contains a Microsoft Entra ID group named HR. The HR group contains two users, User1 and User2. User1 is a member of the HR group, and User2 is not. The HR group is configured for self-service password reset (SSPR). What should you do to enable SSPR for only User1 while keeping it disabled for User2?

Name	Type
storage1	Storage account
container1	Blob container
table1	Storage 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.

Azure Storage Explorer □ □□□□ □□ □□□ □□□ □ □□□?

- A. Task1, Task3, Task4 □
- B. Task2, Task3, Task4 □
- C. Task1, Task2, Task3 □
- D. Task1 □ Task3 □
- E. Task1, Task2, Task3, Task4

Answer: B (LEAVE A REPLY)

NEW QUESTION: 39

□□ □□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

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

App1 □□□ □ □□ West US Azure □□□ □□□□□. App1 □ □□□□ □□□. □□□□ □□□ □□□□□ □□□. □□ □□□□ □□ □□□□ □□□ □□□□ □□□?

- A. □□□□1
- B. □□□3
- C. □□4
- D. □□□2

Answer: D (LEAVE A REPLY)

NEW QUESTION: 40

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

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

Azure Cloud Shell Azure Resource Manager(ARM) PowerShell commands.

```

$adminPassword = Read-Host -Prompt "Enter the administrator password" -AsSecureString

```

- New-AzVm
- New-AzResource
- New-AzTemplateSpec
- New-AzResourceGroupDeployment

- Tag Tag1'
- ResourceGroupName RG1'
- GroupName ManagementGroup1'
- Subscription 9c8bc1cd-7655-4c66-b3ea-a8ee101d8f75

```

- TemplateUri "https://raw.githubusercontent.com/Azure/azure-quickstart-templates/master/101-vm-simple-windows/azuredeploy.json" `
- adminUsername LocalAdministrator -adminPassword $adminPassword -dnsLabelPrefix ContosoVM1

```

Answer:

```

$adminPassword = Read-Host -Prompt "Enter the administrator password" -AsSecureString

```

- New-AzVm
- New-AzResource
- New-AzTemplateSpec
- New-AzResourceGroupDeployment

- Tag Tag1'
- ResourceGroupName RG1'
- GroupName ManagementGroup1'
- Subscription 9c8bc1cd-7655-4c66-b3ea-a8ee101d8f75

```

- TemplateUri "https://raw.githubusercontent.com/Azure/azure-quickstart-templates/master/101-vm-simple-windows/azuredeploy.json" `
- adminUsername LocalAdministrator -adminPassword $adminPassword -dnsLabelPrefix ContosoVM1

```

Explanation:

```
$adminPassword = Read-Host -Prompt "Enter the administrator password" -AsSecureString
```

- New-AzVm
- New-AzResource
- New-AzTemplateSpec
- New-AzResourceGroupDeployment

- Tag Tag1 '
- ResourceGroupName RG1 '
- GroupName ManagementGroup1 '
- Subscription 9c8bc1cd-7655-4c66-b3ea-a8ee101d8f75

```
- TemplateUri "https://raw.githubusercontent.com/Azure/azure-quickstart-templates/master/101-  
- adminUsername LocalAdministrator -adminPassword $adminPassword -dnsLabelPrefix ContosoVM1
```

Reference:

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

NEW QUESTION: 41

VM1 is a virtual machine in a resource group. You need to move VM1 to a different resource group. Which command should you run?

A. `az vm move --resource-group ContosoVM1 --name VM1`

B. `az vm update --resource-group ContosoVM1 --name VM1 --set resource_group_name=ContosoVM2`

C. `az vm update --resource-group ContosoVM1 --name VM1 --set resource_group_name=ContosoVM2 --tags Tag1=ContosoVM1`

D. `az vm update --resource-group ContosoVM1 --name VM1 --set resource_group_name=ContosoVM2 --tags Tag1=ContosoVM2`

- A. `az vm move --resource-group ContosoVM1 --name VM1`
- B. `az vm update --resource-group ContosoVM1 --name VM1 --set resource_group_name=ContosoVM2`

Answer: B (LEAVE A REPLY)

Moving the virtual machine to a different resource group does not change the host that the virtual machine runs on. It only changes the logical grouping of the resources. To move the virtual machine to a different host, you need to redeploy it or use Azure Site Recovery. Then, References: [Move resources to new resource group or subscription] [Redeploy Windows VM to new Azure node] [Use Azure Site Recovery to migrate Azure VMs between Azure regions]

NEW QUESTION: 42

You have an Azure subscription named Subscription1. You need to create a storage account in the subscription. Which Azure service should you use?

Name	Account kind	Azure service that contains data
storage1	Storage	File
storage2	StorageV2 (general purpose v2)	File, Table
storage3	StorageV2 (general purpose v2)	Queue
storage4	BlobStorage	Blob

Azure Import/Export `az storage account create --subscription Subscription1 --name storage1 --kind Storage`

`az storage account create --subscription Subscription1 --name storage2 --kind StorageV2 --sku-name Standard_LRS`

`az storage account create --subscription Subscription1 --name storage3 --kind StorageV2 --sku-name Standard_LRS`

`az storage account create --subscription Subscription1 --name storage4 --kind BlobStorage`

□□□ □□□□ □□□?

- A. □□□□1
- B. □□□□2
- C. □□□□3
- D. □□□□4

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: 43

Azure □□□ □□□□.

Azure Resource Manager □□□□ □□□□ □□□ □□□ □□□ □□□ 50□□ Azure □□ □□□ □□□ □□□□□.

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

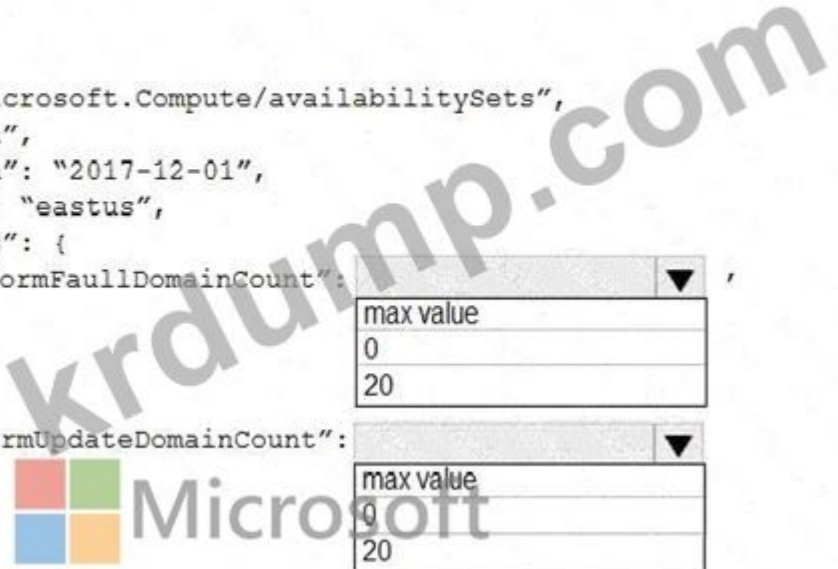
□□□□ □□□ □□□□ □□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

```

"$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": {
      "platformFaultDomainCount": 
      "platformUpdateDomainCount": 
    }
  }
]

```



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": {
        "platformFaultDomainCount": [
          {
            "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: 44

Azure . Windows Server .

Rule1 (DCR) .

Azure Monitor Agent Windows .

ID 1001 .

Rule1 ?

- A. SQL
- B. KQL
- C. XPath

Answer: C (LEAVE A REPLY)

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: 45

☐☐ ☐☐ ☐☐☐ ☐☐ IP ☐☐☐ ☐☐ Azure ☐☐☐ ☐☐☐☐.

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

Bastion1☐☐☐ Azure Bastion Basic SKU ☐☐☐☐ ☐☐☐ ☐☐☐☐☐.

Bastion1☐ ☐☐ IP ☐☐☐ ☐☐☐ ☐ ☐☐☐?

- A. IP1☐
- B. IP1 ☐ IP2☐
- C. IP3, IP4 ☐ IPS☐ ☐☐
- D. IP1, IP2, IP4, IP5☐ ☐☐
- E. IP1, IP2, IP3, IP4 ☐ IPS

Answer: A (LEAVE A REPLY)

When deploying an Azure Bastion host, the network configuration must meet specific requirements depending on the Bastion SKU. For a Bastion Basic SKU, the associated public IP address must fulfill all of the following Azure-defined conditions:

SKU requirement: The Bastion Basic host requires a Standard public IP address.

Tier restriction: The IP must be Regional, not Global, because Bastion operates within a specific Azure region and cannot use global IPs.

Assignment requirement: The IP must use Static allocation, as Bastion requires a reserved, unchanging public IP for consistent access and DNS resolution.

Supported IP version: Azure Bastion supports IPv4 addresses only; IPv6 addresses are not supported.

From the table provided:

IP1 - Standard, Regional, Static, IPv4 # # Meets all requirements.

IP2 - Standard, Global, Static # # Global tier is not supported.

IP3 - Basic, Dynamic # # Bastion requires Standard SKU and Static IP.

IP4 - Basic, Static # # SKU must be Standard.

IP5 - IPv6, Dynamic # # IPv6 and dynamic IPs are unsupported.

Therefore, only IP1 meets the Microsoft Azure Bastion Basic deployment prerequisites.

Reference (Azure Administrator Docs Extract):

Azure Bastion hosts require a Standard, Regional, Static, IPv4 public IP address. Basic or Global IPs are not supported for deployment.

NEW QUESTION: 46

VM3☐ ☐☐☐ ☐☐ ☐☐☐ ☐☐☐☐ ☐☐☐☐ ☐☐☐ ☐☐☐☐☐☐.

- NSG□ □□□ □□□□ □□□.
- □□□□ □□□?
- A. VNet1□ □□□□□
- B. Azure Advisor□ □□ □□ □□
- C. Azure Monitor□ □□ □□
- D. Traffic Manager □□□□ □□ □□ □ □□
- E. Azure Network Watcher□□ IP □□ □□

Answer: E (LEAVE A REPLY)

Scenario: Litware must meet technical requirements including:

Ensure that VM3 can establish outbound connections over TCP port 8080 to the applications servers in the Montreal office.

IP flow verify checks if a packet is allowed or denied to or from a virtual machine. The information consists of direction, protocol, local IP, remote IP, local port, and remote port. If the packet is denied by a security group, the name of the rule that denied the packet is returned. While any source or destination IP can be chosen, IP flow verify helps administrators quickly diagnose connectivity issues from or to the internet and from or to the on-premises environment.

References:

<https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-ip-flow-verify-overview>

AZ-104-KR □□ □□□ □□□□□ □□ DumpTop □□ □□□□ □□□ AZ-104-KR □□! DumpTop □ □□ **AZ-104-KR** □□ □□□ □□□□□□, DumpTop AZ-104-KR □□ □□□ □□□□□ □□□ □□□ □□□□□□□. □□□□ □□□ □□□□ □□ DumpTop AZ-104-KR □□□ □□□□□. <https://www.dumptop.com/Microsoft/AZ-104-KR-dump.html> (454 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 47

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

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

Azure Key Vault□ □□□□ □□□□□□□□ □□□□□ □□□□□ □□□.

app1□ □ □□□□ □□□□□ □□□ □□□□ □□, □□ □ □□□□□ □□ □□□ □□□ □ □□□□□? □□□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□.

ANSWER AREA



Answer:

Answer Area



Explanation:



NEW QUESTION: 48

VNet1 is a virtual network in Azure. It is connected to the Internet.

It contains a virtual machine named VM1. VM1 is connected to the Internet.

Firewalls and virtual networks Private endpoint connections

Save Discard Refresh

Allow access from

All networks Selected networks

Configure network security for your storage accounts. Learn more

Virtual networks

Add existing virtual network Add new virtual network

Virtual Network	Subnet	Address range	Endpoint Status	Resource Group	Subscription
VNET1	1			RG1	Visual Studio Premium with MSDN
	Prod	10.2.0.0/24	Enabled	RG1	Visual Studio Premium with MSDN

Firewall

Add IP ranges to allow access from the internet or your on-premises networks. Learn more.

Add your client IP address (51.145.137.40)

Address range

IP address or CIDR

Resource instances

Specify resource instances that will have access to your storage account based on their system-assigned managed identity. Rules created by other tenants can only be modified by the creator.

Resource type

Select a resource type

Instance name

Select one or more instances

Exceptions

- Allow trusted Microsoft services to access this storage account
- Allow read access to storage logging from any network
- Allow read access to storage metrics from any network

Network Routing

Determine how you would like to route your traffic as it travels from its source to an Azure endpoint. Microsoft routing is recommended for most customers.

Routing preference *

Microsoft network routing Internet routing

Publish route-specific endpoints

- Microsoft network routing
- Internet routing

0000 000 000 0000 0 000 0000 00 000 00000 0000 000 00000.
00: 00 000 10000.

Answer Area

The virtual machines on the 10.2.9.0/24 subnet will have network connectivity to the file shares in the storage account **[answer choice]**.

never
always
during a backup
never

Azure Backup will be able to back up the unmanaged hard disks of the virtual machines in the storage account **[answer choice]**.

never
always
during a backup
never

Answer:

Answer Area

The virtual machines on the 10.2.9.0/24 subnet will have network connectivity to the file shares in the storage account **[answer choice]**.

Azure Backup will be able to back up the unmanaged hard disks of the virtual machines in the storage account **[answer choice]**.

never
always
during a backup
never

never
always
during a backup
never

Explanation:

Never

Never

NEW QUESTION: 49

Subscription1 is an Azure subscription.

contosostorage is an Azure Storage account in Subscription1.

contosostorage has a file share named UNC. UNC is a file share that is not managed by Azure Backup.

contosostorage has a file share named data. data is a file share that is managed by Azure Backup.

contosostorage has a file share named 1. 1 is a file share that is managed by Azure Backup.

Values

- blob
- blob.core.windows.net
- contosostorage
- data
- file
- file.core.windows.net
- portal.azure.com
- subscription1

Answer Area

\\ . \

Answer:

Values

- blob
- blob.core.windows.net
- contosostorage
- data
- file
- file.core.windows.net
- portal.azure.com
- subscription1

Answer Area

\\ contosostorage . file.core.windows.net \ data

Explanation:

Values

- blob
- blob.core.windows.net
- contosostorage
- data
- file
- file.core.windows.net
- portal.azure.com
- subscription1

Answer Area

\\contosostorage.file.core.windows.net\data

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: 50

App1 Azure NSG NSG NSG NSG NSG.

NSG NSG (NSG) NSG NSG NSG.

NSG NSG App1 NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG.

NSG NSG NSG?

- A. NSG NSG 443 NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG.
- B. NSG NSG 443 NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG.
- C. NSG NSG 443 NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG.
- D. NSG NSG 443 NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG NSG.

Answer: C (LEAVE A REPLY)

As App1 is public-facing we need an incoming security rule, related to the access of the web servers.

Scenario: You have a public-facing application named App1. App1 is comprised of the following three tiers: a SQL database, a web front end, and a processing middle tier.

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

Topic 3, Contoso Ltd (Consulting Company) Case study

This is a case study. Case studies are not timed separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other questions in this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next section of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Overview

General Overview

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

Environment

Existing Environment

Contoso has an Azure subscription named Sub1 that is linked to an Azure Active Directory (Azure AD) tenant. The network contains an on-premises Active Directory domain that syncs to the Azure AD tenant.

The Azure AD tenant contains the users shown in the following table.

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

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

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

User1 manages the resources in RG1. User4 manages the resources in RG2.

Sub1 contains virtual machines that run Windows Server 2019 as shown in the following table

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

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

Sub1 contains the storage accounts shown in the following table.

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

Requirements

Planned Changes

Contoso plans to implement the following changes:

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

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

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

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

Associate NSG1 to the network interface of VM1.

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

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

Associate NSG2 to VNET1/Subnet2.

Technical Requirements

Contoso must meet the following technical requirements:

Create container1 and share1.

Use the principle of least privilege.

Create an Azure AD security group named Group4.

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

Trigger an alert if VM1 or VM2 has less than 20 GB of free space on volume C.

Enable User1 to create Azure policy definitions and User2 to assign Azure policies to RG1.

Create an internal Basic Azure Load Balancer named LB1 and connect the load balancer to VNET1/Subnet1 Enable flow logging for IP traffic from VM5 and retain the flow logs for a period of eight months.

Whenever possible, grant Group4 Azure role-based access control (Azure RBAC) read-only permissions to the Azure file shares.

NEW QUESTION: 51

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

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

Azure Monitor DCRI can collect data from the following resources.

DCRI can collect data from the following resources, which are not supported by DCRI? Select all that apply.

Options: VM1, storage1, Workspace1, DB1.

Answer Area  Microsoft

Data sources:

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

Destinations:

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

Answer:

Answer Area

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: 52

Adatum is a Microsoft Azure Active Directory (Azure AD) tenant. The tenant is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources.

Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources.

Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources.

Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources.

Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources.

Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources. Adatum is used to manage the access of users to Azure resources.

A.

B.

Answer: [SHOW ANSWER](#)

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

[Azure Logic Apps permissions and access control]

NEW QUESTION: 53

Sub1 Azure, User1 User2

RG1

RG2

VM1

VNet1

VNet2

User1 Sub1 User2 RG1

Azure Resource Manager(ARM)

- * CPU: 8
* RAM: 64GB
* VNet2
* RG1
* Template2
* CPU: 16
* RAM: 192GB
* VNet1
* RG1

Policy1 128GB RAM

- * RG1
*
*

Policy1

Statements	Yes	No
VM1 is within the scope of Policy1.	<input type="radio"/>	<input type="radio"/>
User1 can deploy a virtual machine by using Template2.	<input type="radio"/>	<input type="radio"/>
User2 can deploy a virtual machine by using Template1.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
VM1 is within the scope of Policy1.	<input type="radio"/>	<input checked="" type="radio"/>
User1 can deploy a virtual machine by using Template2.	<input type="radio"/>	<input checked="" type="radio"/>
User2 can deploy a virtual machine by using Template1.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

No, No, No

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

Policy1 is assigned only to RG1 , with no exclusions. VM1 is in RG2 , so it is outside the assignment scope and is not within the scope of Policy1. Microsoft's Azure Policy guidance states that the policy rule determines which resources in the assignment scope are evaluated.

User1 has the Owner role at the subscription scope, which normally grants full management access, including access to resources under RG1. However, Template2 deploys a VM with 192 GB RAM to RG1. Because Policy1 denies deployment of virtual machines with more than 128 GB of RAM, the request is blocked even for an Owner. Microsoft defines the Deny effect as preventing a noncompliant resource request and failing the request.

User2 has Contributor only on RG1 . Template1 deploys the VM into RG1 and uses only 64 GB RAM, so Policy1 does not deny it. However, Template1 connects the VM to VNet2 , which is in RG2 . User2 has no permissions on RG2 or VNet2, so the deployment cannot complete. Study Guide reference: AZ-104 Manage Azure identities and governance , including Azure RBAC, role scope, and Azure Policy.

NEW QUESTION: 54

□□□ □□□ Azure□ □□□ □□□.

□□□ □□ □□□?

A. □□ □□□ □□(SAS)□ □□□□□. □□□□□ □□□ □□ □□ □□□□ □□□□ □□□ □□□□□.

B. Azure Import/Export □□□□ □□□□□.

C. □□□ □□ □□□□□. □□□□□ □□□ □□ □□ □□□□ □□□□ □□□ □□□□□.

D. Azure Storage Explorer□ □□□□ □□□ □□□□□.

Answer: D (LEAVE A REPLY)

Azure Storage Explorer is a free tool from Microsoft that allows you to work with Azure Storage data on Windows, macOS, and Linux. You can use it to upload and download data from Azure blob storage.

Scenario:

Planned Changes include: move the existing product blueprint files to Azure Blob storage.

Technical Requirements include: Copy the blueprint files to Azure over the Internet.

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

NEW QUESTION: 55

Azure Network Watcher is a service that provides network monitoring and troubleshooting for Azure virtual networks.

Scenario 1: A virtual network is connected to the Internet. A virtual machine in the virtual network is unable to access the Internet. Scenario 2: A virtual network is connected to the Internet. A virtual machine in the virtual network is unable to access the Internet.

Task 1: Troubleshoot the connectivity issue.

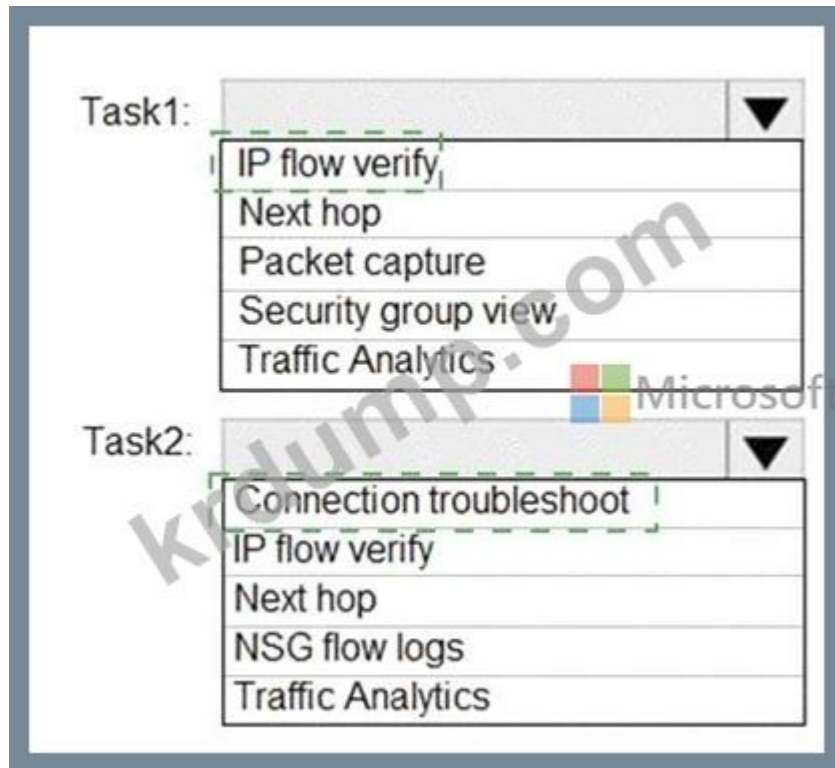
Task1: ▼

IP flow verify
Next hop
Packet capture
Security group view
Traffic Analytics

Task2: ▼

Connection troubleshoot
IP flow verify
Next hop
NSG flow logs
Traffic Analytics

Answer:



Explanation:

Task 1: IP flow verify

The IP flow verify capability enables you to specify a source and destination IPv4 address, port, protocol (TCP or UDP), and traffic direction (inbound or outbound). IP flow verify then tests the communication and informs you if the connection succeeds or fails. If the connection fails, IP flow verify tells you which security rule allowed or denied the communication, so that you can resolve the problem.

Task 2: Connection troubleshoot

The connection troubleshoot capability enables you to test a connection between a VM and another VM, an FQDN, a URI, or an IPv4 address. The test returns similar information returned when using the connection monitor capability, but tests the connection at a point in time, rather than monitoring it over time.

Reference:

<https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-monitoring-overview>

<https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-ip-flow-verify-overview>

<https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-connectivity-overview>

NEW QUESTION: 56

□□□□□ □□□□□ Active Directory □□□ □□□(AD DS) □□□□ □□□□ □□□□.

□□ □□□□□ □□ □□ □□□ ID□ □□□□ □□□□.

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

AD DS □□□□□ Microsoft Entra □□□□□ OU1□ □□□□□ □□ Microsoft Entra Connect □□□□ □□□□ Azure □□□ □□□□.

Microsoft Entra □□□□□ User3□□□□ □□□□ □□ □□□□ □□□□ □□□□.

Answer Area



Statements	Yes	No
User1 can access content in share1.	<input type="radio"/>	<input type="radio"/>
User2 can access content in share1.	<input type="radio"/>	<input type="radio"/>
User3 can access content in share1.	<input type="radio"/>	<input type="radio"/>

Answer:



Explanation:

NO

YES

NO

In this scenario, Microsoft Entra Connect (formerly Azure AD Connect) is configured to synchronize only the Organizational Unit (OU1) from the on-premises Active Directory Domain Services (AD DS) environment to the Microsoft Entra tenant (Azure AD).

The synchronization scope determines which objects (users, groups, devices) are replicated to Azure AD. Any objects outside the selected OUs are not synchronized and therefore do not exist in Azure AD.

Here's what happens to each user:

* User1 (in OU2) - Because OU2 is not included in the synchronization scope, User1 is not synchronized to Azure AD. Therefore, User1 does not have an account in Microsoft Entra ID and cannot authenticate to any cloud-based resource such as Share1 (which is assumed to be an Azure file share or Microsoft 365 resource).

* Result: No access.

* User2 (in OU1) - Since OU1 is synchronized using Microsoft Entra Connect, User2 exists in Azure AD. User2's identity is recognized by Azure AD and can be used to access cloud-based resources (like Share1) if permissions are assigned appropriately (either directly or via a group such as Group1).

* Result: Yes, access is possible.

* User3 (cloud-only account) - While User3 exists in Azure AD, there's no indication that this user has been assigned access to Share1. Because access is granted through synchronization (from AD DS) and group membership in synchronized objects (like Group1), User3 would not automatically inherit those permissions.

* Result: No access.

Additionally, Group1, which contains User1 and resides in OU1, will be synchronized. However, since User1 is in OU2, User1's membership in Group1 is not synchronized, because non-synchronized objects cannot appear in synchronized group memberships in Azure AD.

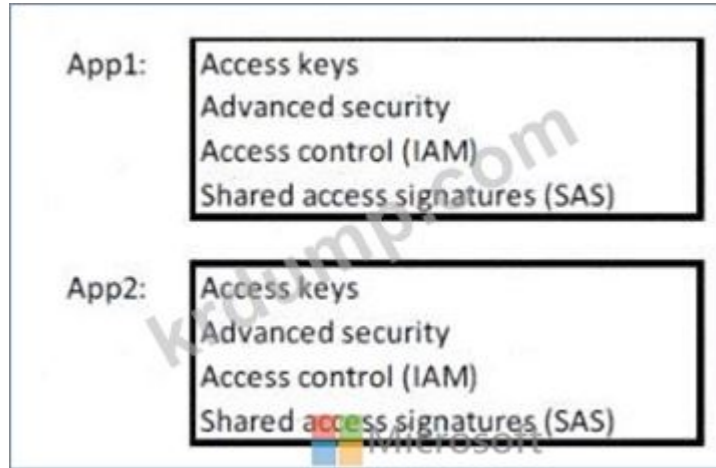
This behavior is defined by Microsoft's Azure AD Connect synchronization rules, which state:

"Objects outside of the configured synchronization scope are not synchronized to Azure AD. Group memberships including users from outside the synchronization scope will not include those users in Azure AD."

NEW QUESTION: 57

storage1□□□ Azure Storage □□□ □□□□.

app1 Azure App Service Azure App2 ID.
 App1 App2 30 storage1 blob storage1?



Answer:



Explanation:

Box 1: Access Control (IAM)

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

Box 2: Shared access signatures (SAS)

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

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

Reference:

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

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

NEW QUESTION: 58

Azure App Service Azure App2 ID.

Azure App1 App2 30 storage1 blob storage1?

Azure App1 App2 30 storage1 blob storage1?

App1 App2 30 storage1 blob storage1?

App1 App2 30 storage1 blob storage1?

- A. `docker tag <image-name> <registry-name>.azurecr.io/<image-name>:<tag>`
- B. `docker tag <image-name> <registry-name>:<tag>`
- C. `docker tag <image-name> <registry-name>`
- D. `docker tag <image-name> <registry-name>:<tag>`

Answer: A (LEAVE A REPLY)

To push a container image to Azure Container Registry (ACR) using the Azure CLI, the image must first be tagged with the registry's login server name. This is a mandatory step in the Docker workflow.

After signing in to the registry using `az acr login`, the local container image must be tagged in the following format:

`< registry-name > .azurecr.io/ < image-name > : < tag >`

Azure documentation specifies that Docker determines the target registry based on the image tag. If the image is not tagged with the registry's login server name, Docker will not know where to push the image.

Listing images, deploying container groups, or configuring YAML files are optional or later steps and are not required to push an image to ACR.

Final Answer: A. Tag a container image with the name of the container registry 's login server

NEW QUESTION: 59

VNet1 is a virtual network in Azure. VM1 is a virtual machine in VNet1. VM1 is currently configured with the following settings:

IP address: 10.0.0.0/24

Availability Set: AVSet

Network Security Group (NSG): None

Public IP address: 10.0.0.4 (Static)

Public IP address: 40.90.219.6 (Dynamic)

Load balancer: slb1 (Backend pool: VM1)

VM1 is currently in a state that prevents it from being connected to by slb1.

What must you do to ensure that VM1 can be connected to by slb1?

Options:

- A. Create a new NSG and assign it to VM1.
- B. Remove the public IP address from VM1.
- C. Change the private IP address of VM1 to static.
- D. Create a new NSG and assign it to slb1.

Before you create a backend pool on slb1, you must:

- Create and assign an NSG to VM1
- Remove the public IP address from VM1
- Change the private IP address of VM1 to static

Before you can connect to VM1 from slb1, you must:

- Create and configure an NSG
- Remove the public IP address from VM1
- Change the private IP address of VM1 to static

Answer:

Before you create a backend pool on slb1, you must:

- Create and assign an NSG to VM1
- Remove the public IP address from VM1
- Change the private IP address of VM1 to static

Before you can connect to VM1 from slb1, you must:

- Create and configure an NSG
- Remove the public IP address from VM1
- Change the private IP address of VM1 to static

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 whether 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.

Answer Area

Before you create a backend pool on slb1, you must:

- Create and assign an NSG to VM1
- Remove the public IP address from VM1
- Change the private IP address of VM1 to static

Before you can connect to VM1 from slb1, you must:

- Create and configure an NSG
- Remove the public IP address from VM1
- Change the private IP address of VM1 to static

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: 60

WebApp1 is an Azure App Service web application. You need to configure DNS for WebApp1.

* WebApp1 must use the app.contoso.com domain.

* WebApp1 must be accessible from the Internet.

* The DNS records must be created automatically.

Which configuration should you use for WebApp1? Select the correct answer.

Options: A. Standard B. Basic C. Free D. Shared

Answer Area

Pricing plan: Standard

- Basic
- Free
- Shared
- Standard

Record type: TXT

- A
- AAAA
- PTR
- TXT

Answer:
Answer Area



Pricing plan: Standard

- Basic
- Free
- Shared
- Standard

Record type: TXT

- A
- AAAA
- PTR
- TXT

Explanation:



When configuring an Azure App Service app (WebApp1), several hosting plans determine scaling capabilities and supported features such as custom domains, SSL, and auto-scaling.

According to Microsoft Azure Administrator Documentation:

* Custom Domain Verification: To map a custom domain (like app.contoso.com) to an Azure Web App, Azure must verify that you own the domain. Microsoft documentation specifies:

"To verify a custom domain, create a TXT record in your DNS zone. The TXT record contains a verification token that Azure uses to confirm ownership of the domain before binding it to your App Service." (Source: Azure App Service - Map custom domain) While an A record or CNAME record is eventually used to direct traffic to the app, the TXT record is used solely for domain verification.

* Scaling Requirement (Up to 8 Instances): The Free and Shared App Service plans have significant limitations:

* Free / Shared: No scaling (only 1 instance, no custom domain SSL support).

* Basic: Supports up to 3 instances.

* Standard: Supports up to 10 instances, includes auto-scaling, and supports custom domains and SSL.

* Premium: Adds advanced scaling and isolation but is more expensive and unnecessary here.

Microsoft documentation states:

"The Standard pricing tier supports auto-scaling up to 10 instances and custom domains, offering a balance between cost and capability." (Source: Azure App Service Plan Tiers Overview) Since the requirement includes automatic scaling up to eight instances and custom domain verification, the Standard plan is the minimum suitable tier - satisfying both performance and cost-efficiency conditions.

Final Verified Answer:

* Pricing plan: Standard

* Record type: TXT

Justification Summary (from Microsoft Documentation):

* Custom domain verification # Requires a TXT record.

* Auto-scale up to 8 instances # Requires at least the Standard plan.

* Minimize cost and effort # Standard plan is the optimal balance.

NEW QUESTION: 61

VM1 is a virtual machine in an Azure virtual network. VM1 is configured for Azure Site Recovery. What is the minimum number of VMs that must be configured for Azure Site Recovery to protect VM1?

- A. Azure Automation
- B. Azure Site Recovery
- C. VM1
- D. Azure Site Recovery

Answer: [\(SHOW ANSWER\)](#)

AZ-104-KR ☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop ☐☐ ☐☐☐☐ ☐☐☐ AZ-104-KR ☐☐! DumpTop ☐ ☐☐ **AZ-104-KR** ☐☐ ☐☐☐ ☐☐☐☐☐☐, DumpTop AZ-104-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐ ☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐. ☐☐☐☐☐ ☐☐☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop AZ-104-KR ☐☐☐☐ ☐☐☐☐☐☐. <https://www.dumptop.com/Microsoft/AZ-104-KR-dump.html> (454 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 62

☐☐ ☐☐ ☐☐☐☐ ☐☐☐☐☐ ☐☐☐ Azure ☐☐☐☐ ☐☐☐☐☐.

☐☐☐☐ ☐☐☐☐☐☐☐☐☐☐.

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

RG1 | Access control (IAM) ...
Resource group

Search (Ctrl+/) « + Add ↓ Download role assignments ≡ Edit columns ↻ Refresh | ✕ Remove | 🗨 Got feedback?

Overview
Activity log
Access control (IAM)
Tags
Resource visualizer
Events

Settings
Deployments
Security
Policies
Properties
Locks

Check access Role assignments Roles Deny assignments Classic administrators

Number of role assignments for this subscription ⓘ
2 2000

Search by name or email Type: All Role: All Scope: All scopes Group by: Role

2 items (1 Users, 1 Groups)

<input type="checkbox"/>	Name	Type	Role	Scope	Condition
▼	Owner				
<input type="checkbox"/>	Group1	Group	Owner ⓘ	This resource	None
<input type="checkbox"/>	prvi prvi...	User	Owner ⓘ	Subscription (Inherited)	None

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

Answer Area

Statements	Yes	No
You can assign User2 the Owner role for RG1 by adding Group2 as a member of Group1.	<input type="radio"/>	<input type="radio"/>
You can assign User3 the Owner role for RG1 by adding Group3 as a member of Group1.	<input type="radio"/>	<input type="radio"/>
You can assign User3 the Owner role for RG1 by assigning the Owner role to Group3 for RG1.	<input type="radio"/>	<input type="radio"/>

Answer:

ANSWER AREA

Statements

You can assign User2 the Owner role for RG1 by adding Group2 as a member of Group1.

Yes

No

You can assign User3 the Owner role for RG1 by adding Group3 as a member of Group1.



Yes

No

You can assign User3 the Owner role for RG1 by assigning the Owner role to Group3 for RG1.

Explanation:

Answer Area

Statements

You can assign User2 the Owner role for RG1 by adding Group2 as a member of Group1.

You can assign User3 the Owner role for RG1 by adding Group3 as a member of Group1.

You can assign User3 the Owner role for RG1 by assigning the Owner role to Group3 for RG1.

Yes No

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

" Group nesting isn ' t supported. A group can ' t be added as a member of a role-assignable group. " For the second question:

<https://learn.microsoft.com/en-us/azure/active-directory/fundamentals/how-to-manage-groups#add-or-remove-a-group-from-another-group>

" We currently don ' t support:

Adding Microsoft 365 groups to Security groups or other Microsoft 365 groups.

"

For the third question, although it appears truncated in the screenshot (ending with " for... ") there is a reference about Microsoft 365 groups support for roles assignment here:

<https://learn.microsoft.com/en-us/azure/active-directory/roles/groups-concept#how-role-assignments-to-groups-work>

" To assign a role to a group, you must create a new security or Microsoft 365 group with the isAssignableToRole property set to true. "

NEW QUESTION: 63

contoso.com □□□ Microsoft Entra □□□□ □□□□.

abrikam.com □□□ □□ □□□□ □□□□□.

abrikam.com □ □□□□ contoso.com □□□□ □□□□□ □□□.

□□□□ abrikam.com □□□□□□ □□ □ □□□ □□ □□□.

Microsoft Entra □□ □□□□ □□□ □□ □□□?

A. □□ □□ □□□□ □□□ □□□ □□□ □□ □□□ □□□□□.

B. □□ □□□ □□□ □□□□ □□□ □□ □□□ □□□□□.

C. □□ □□ □□□□ □□ □□ □□□ □□□□□.

D. □□ □□□ □□□ □□□□ Microsoft □□□□ □□□ □□□□□.

Answer: C (LEAVE A REPLY)

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

Within the External collaboration settings, administrators can:

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

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

The other options do not meet the requirement:

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

Final Verified Answer:

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

NEW QUESTION: 64

Subscription1 is an Azure subscription. Subscription1 contains VM1. VM1 is connected to the Internet.

VM1 is connected to the Internet through a public IP address.

VM1 is connected to the Internet through a public IP address.

Network Interface: vm1441 Effective security rules Topology

Virtual network/subnet: VNET1/default NIC Public IP: 52.160.123.200 NIC Private IP: 10.0.6.4 Accelerated networking: Disabled

Inbound port rules Outbound port rules Application security groups Load balancing

Network security group VM1-nsg (attached to network interface: vm1441)
Impacts 0 subnets, 1 network interfaces [Add inbound port rule](#)

Priority	Name	Port	Protocol	Source	Destination	Action
100	Rule2	50-60	Any	Any	Any	Deny
300	RDP	3389	TCP	Any	Any	Allow
400	Rule1	50-500	Any	Any	Any	Allow
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	AllowAzureLoadBalancerInBound	Any	Any	AzureLoadBalancer	Any	Allow

VM1 is connected to the Internet through a public IP address.

VM1 is connected to the Internet through a public IP address.

Answer Area

Internet users [answer choice].

- can connect to only the web server on VM1
- can connect to only the DNS server on VM1
- can connect to only the web server on VM1
- can connect to the web server and the DNS server on VM1
- cannot connect to the web server and the DNS server on VM1

If you delete Rule2, Internet users [answer choice].

- can connect to the web server and the DNS server on VM1
- can connect to only the DNS server on VM1
- can connect to only the web server on VM1
- can connect to the web server and the DNS server on VM1
- cannot connect to the web server and the DNS server on VM1

Answer:
Answer Area

Internet users [answer choice].

- can connect to only the web server on VM1
- can connect to only the DNS server on VM1
- can connect to only the web server on VM1
- can connect to the web server and the DNS server on VM1
- cannot connect to the web server and the DNS server on VM1

If you delete Rule2, Internet users [answer choice].

- can connect to the web server and the DNS server on VM1
- can connect to only the DNS server on VM1
- can connect to only the web server on VM1
- can connect to the web server and the DNS server on VM1
- cannot connect to the web server and the DNS server on VM1

Explanation:

Answer Area

Internet users [answer choice].

can connect to only the web server on VM1

If you delete Rule2, Internet users [answer choice].

can connect to the web server and the DNS server on VM1

In this scenario, the network security group (NSG) vm1441-nsg controls inbound traffic to VM1. Azure NSGs operate using priority-based rules, where the lower number indicates a higher priority. Rules are processed in ascending order until a match is found.

From the NSG configuration in the exhibit:

- * Rule2 (Priority 100): Deny traffic on ports 50-60.
- * Rule "RDP" (Priority 300): Allow TCP 3389 (RDP).
- * Rule1 (Priority 400): Allow ports 50-500.
- * Default rules (priority 65000-65500) permit traffic within the VNet and Azure Load Balancer, and deny all others.

Azure NSGs apply first-match logic - once a rule matches, no subsequent rules are processed. The lower- numbered Rule2 (priority 100) takes precedence over Rule1 (priority 400), meaning ports 50-60 are denied before Rule1 can allow them.

Given that:

* The web server typically runs on port 80 (HTTP) or 443 (HTTPS), which are not affected by Rule2 - they remain allowed by default or open firewall configuration.

* The DNS server uses UDP/TCP port 53, which falls within the denied range (50-60) under Rule2.

Therefore:

* While Rule2 exists, Internet users can only access the web server (port 80/443).

* Once Rule2 is deleted, the deny restriction for ports 50-60 is removed, meaning both the web server (80/443) and DNS server (53) are accessible.

This aligns directly with Microsoft Azure documentation for NSGs (Microsoft Learn: "Network security groups filter network traffic to and from Azure resources by using 5-tuple rules (source, destination, port, protocol, and direction). Rules are processed in priority order, and the first match wins.").

NEW QUESTION: 65

contoso.com is an Azure Directory (Azure AD) tenant. It contains three users: SecAdmin1, BillAdmin1, and User1.

SecAdmin1 is a Security administrator. BillAdmin1 is a Billing administrator. User1 is a Reports reader.

Name	Role
SecAdmin1	Security administrator
BillAdmin1	Billing administrator
User1	Reports reader

The tenant contains three roles: Security administrator, Billing administrator, and Reports reader. Each role is assigned to a user.

Statement 1: SecAdmin1 can reset the password of User1.

Statement 2: BillAdmin1 can reset the password of SecAdmin1.

Statement 3: User1 can reset the password of SecAdmin1.

Statement 4: User1 can reset the password of BillAdmin1.

Statement 5: User1 can reset the password of User1.

Statement 6: SecAdmin1 can reset the password of BillAdmin1.

Statement 7: SecAdmin1 can reset the password of SecAdmin1.

Statement 8: BillAdmin1 can reset the password of User1.

Statement 9: User1 can reset the password of BillAdmin1.

Statement 10: User1 can reset the password of User1.

Answer Area



Microsoft

Statements

Yes No

SecAdmin1 must answer the following question if he wants to reset his password:
In what city was your first job?

Yes No

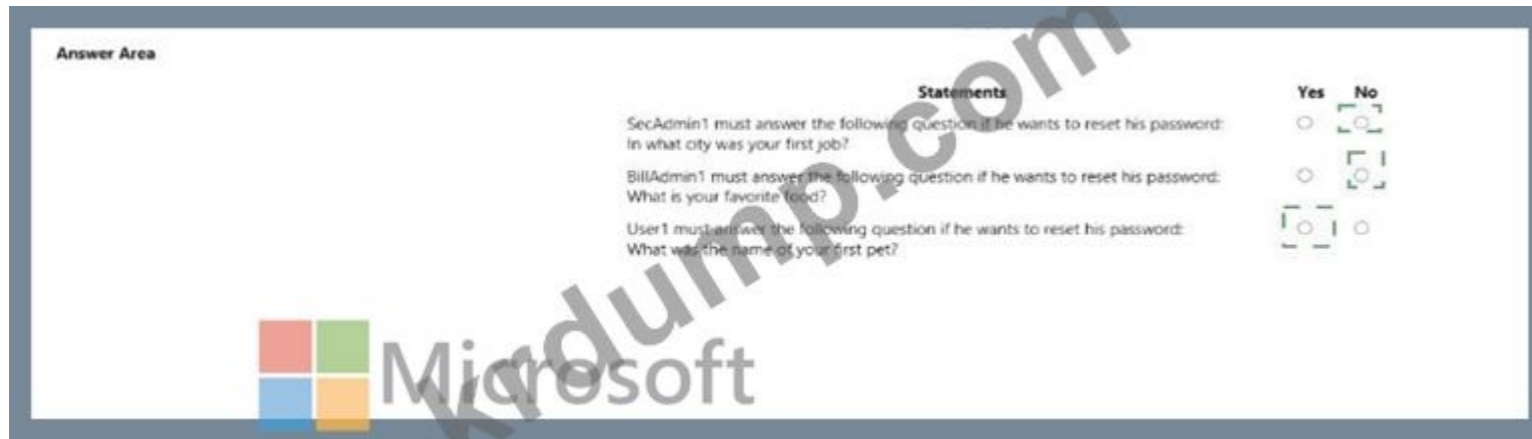
BillAdmin1 must answer the following question if he wants to reset his password:
What is your favorite food?

Yes No

User1 must answer the following question if he wants to reset his password:
What was the name of your first pet?

Yes No

Answer:



Explanation:

No, No, Yes

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

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

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

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

NEW QUESTION: 66

Company has two subscriptions, Sub1 and Sub2, in Azure. The following table shows the configuration of virtual machines (VMs) in each subscription.

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

VM1 and VM2 are connected to the Internet via a virtual private network (VPN) gateway in Sub1. The VPN gateway is configured to allow traffic from the Internet to the VMs.

- A. Azure VPN gateway
- B. Azure firewall
- C. Azure user-defined routes (UDR)
- D. Azure network security groups (NSG)
- E. Azure network virtual appliances (NVA)

Answer: D (LEAVE A REPLY)

Azure Virtual Network (VNet) Peering is the Microsoft-recommended solution to enable seamless communication between two virtual networks, either within the same subscription or across different subscriptions, with minimal cost and administrative overhead.

According to the Microsoft Azure Administrator documentation, VNet Peering connects two virtual networks and allows resources in either VNet to communicate with each other using private IP addresses, just as if they were part of the same network. Traffic between peered VNets remains on the Microsoft backbone network, ensuring low latency, high bandwidth, and no data exposure to the public internet.

There are two types of peering supported:

- * VNet Peering (intra-region): For virtual networks in the same Azure region.
- * Global VNet Peering: For virtual networks in different Azure regions (for example, East US and West US).

In this scenario, VNet1 (East US, Sub1) and VNet2 (West US, Sub2) are in different regions and different subscriptions. Therefore, Global VNet Peering is the appropriate configuration. Peering can be established across subscriptions, provided that the subscriptions are associated with the same Azure Active Directory tenant or proper permissions exist between tenants.

Unlike VPN gateways or network virtual appliances, VNet peering does not require additional infrastructure or incur data transfer gateway costs, making it the lowest-cost and least administratively complex option.

Once configured, communication between VM1 in VNet1 and VM2 in VNet2 occurs over Microsoft's private backbone network, without public IPs or tunneling.

NEW QUESTION: 67

RG1 is a resource group in Azure. It contains the following resources:
 storage1 is a storage account in RG1.
 File1 is a file share in storage1.
 storage1 is in the same region as RG1. File1 is in the same region as storage1.
 What is the scope of File1?

- A. location
- B. kind
- C. sku
- D. scope

Answer: [\(SHOW ANSWER\)](#)

NEW QUESTION: 68

VM1 and VM2 are virtual machines in Azure. They are configured as follows:

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

VM1 and VM2 have private IP addresses. VM1 and VM2 are Windows Server 2019. Subnet1 and Subnet2 are in VNET1. NSG1 and NSG2 are network security groups (NSG) in VNET1. NSG1 is associated with Subnet1. NSG2 is associated with Subnet2.

- * Subnet1: 100
- * Subnet2: 101
- * VNET1: 3389

- * □□□□ : TCP
- * □□ : Any
- * □□□ : □□□
- * □□: □□

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

Answer Area

Statements	Yes	No
From the internet, you can connect to VM1 by using Remote Desktop.	<input type="radio"/>	<input type="radio"/>
From the internet, you can connect to VM2 by using Remote Desktop.	<input type="radio"/>	<input type="radio"/>
From VM1, you can connect to VM2 by using Remote Desktop.	<input type="radio"/>	<input type="radio"/>

Answer:



Answer Area

Statements	Yes	No
From the internet, you can connect to VM1 by using Remote Desktop.	<input type="radio"/>	<input checked="" type="radio"/>
From the internet, you can connect to VM2 by using Remote Desktop.	<input checked="" type="radio"/>	<input type="radio"/>
From VM1, you can connect to VM2 by using Remote Desktop.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
From the internet, you can connect to VM1 by using Remote Desktop.	<input type="radio"/>	<input checked="" type="radio"/>
From the internet, you can connect to VM2 by using Remote Desktop.	<input checked="" type="radio"/>	<input type="radio"/>
From VM1, you can connect to VM2 by using Remote Desktop.	<input checked="" type="radio"/>	<input type="radio"/>

In Azure, Network Security Groups (NSGs) control inbound and outbound traffic to network interfaces (NICs), subnets, and virtual machines (VMs) using rules based on priority and direction. According to the Microsoft Azure Administrator Guide and Azure Networking documentation, the following principles apply:

Default NSG rules:

By default, an NSG denies all inbound traffic from the Internet except traffic originating from the same Virtual Network (VNet).

NSG allows all outbound traffic to the Internet.

Default inbound rules include:

Allow VNet inbound (priority 65000)

Allow Azure Load Balancer inbound (priority 65001)

Deny all inbound (priority 65500)

NSG associations:

NSGs can be associated either with a subnet or an individual network interface (NIC).

When both are applied, the NIC-level NSG takes precedence for inbound/outbound traffic.

If no explicit allow rule exists, default deny applies.

Analysis of Each Statement

1## From the internet, connect to VM1 by using RDP:

VM1's subnet (Subnet1) is associated with NSG1, which has only the default rules.

The default rules deny all inbound traffic from the Internet, including port 3389 (RDP).

Therefore, RDP from the Internet is blocked.

The Answer: # No

2## From the internet, connect to VM2 by using RDP:

VM2's NIC is associated with NSG2, which includes a custom allow rule (priority 100) permitting TCP traffic on port 3389 from any source to any destination.

This rule overrides the default deny rule.

Thus, RDP from the Internet is allowed.

The Answer: # Yes

3## From VM1, connect to VM2 by using RDP:

Both VMs reside in the same VNet (VNET1) but different subnets (Subnet1 and Subnet2).

The default NSG rule "Allow VNet Inbound" allows traffic between subnets within the same virtual network.

Therefore, VM1 can connect to VM2 via RDP (port 3389).

The Answer: # Yes

NEW QUESTION: 69

□□ □□ Azure □□□ □□□□, □ □□□□ □□ □□ □□ □□□□ □□□□.

VM1□□□ □□ □□□ □□□□ □□ □□ □ □□ □□□□□ □□□□□.

Azure Monitor□ □□ □□ □□□ □□ □□□□ □□□.

A. □□□ □□ □□(DCR)□ □□□□□.

B. □□ □□□ □□□□□.

C. □□ □□□ □□□□.

D. □□ □□□ □□□□□.

Answer: [\(SHOW ANSWER\)](#)

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

The correct first action is to create a data collection rule (DCR) . Azure Monitor can collect host-level platform metrics automatically, but metrics from inside the guest operating system require an agent-based collection path. Microsoft states that Azure Monitor Agent collects monitoring data from the guest operating system of Azure and hybrid VMs, and that it collects data according to data collection rules, which define what data is collected, processed, and sent.

Disk read and write activity inside the OS is represented through performance counters. Microsoft's VM monitoring guidance confirms that guest OS metrics must be collected through agents, and DCRs define performance counters and destinations such as Log Analytics workspaces. A workbook is useful later for consolidated visualization across subscriptions, because workbooks combine metrics, log queries, and multiple data sources into interactive reports, but it does not initiate guest OS metric collection. An alert rule detects conditions after data exists, and diagnostic settings mainly route platform logs or supported platform metrics rather than collecting guest OS counters. This aligns with AZ-104 Monitor resources in Azure, especially configuring Azure Monitor log settings, interpreting metrics, querying logs, and monitoring virtual machines by using Azure Monitor.

NEW QUESTION: 70

Azure .

.

Azure .

?

- A. Azure
- B. Azure Container Instances
- C. Azure Container Apps Azure App Service
- D. Azure Container Instances Azure App Service
- E. Azure , Azure Azure App Service

Answer: C (LEAVE A REPLY)

According to Microsoft Azure Administrator (AZ-104) study guides and Microsoft's official documentation on container hosting options, the ability to automatically scale containerized workloads depends on the underlying platform's orchestration and scaling capabilities.

Azure Container Apps (ACA) is a fully managed service that runs microservices and containers without requiring orchestration infrastructure such as Kubernetes. It natively supports automatic scaling (autoscaling) through KEDA (Kubernetes Event-Driven Autoscaler) integration. Autoscaling in ACA can be triggered by various metrics-such as HTTP request rate, CPU utilization, memory usage, or custom event sources like Azure Service Bus or Event Hub. Administrators can configure minimum and maximum replicas, and Azure dynamically adjusts container instances to match load demand. Azure App Service can also host containerized web applications (via custom Docker images) and includes built-in autoscaling capabilities. App Service Plans support both manual and automatic scaling based on metrics like CPU percentage, memory percentage, and HTTP queue length. This scaling can occur horizontally (by adding instances) or vertically (by changing the pricing tier). In contrast, Azure Container Instances (ACI) provides simple, isolated container execution and does not support automatic scaling. You must manually script or orchestrate scaling logic via Azure Automation or external schedulers.

Therefore, per official Microsoft study content and Azure documentation, only Azure Container Apps and Azure App Service offer native, built-in autoscaling for containers.

Final Verified Answer: # C. Azure Container Apps or Azure App Service only

NEW QUESTION: 71

Subscription1 Azure . Subscription1 VM1 VM2 Azure . VM1 VM2 Windows Server 2016 .

VM1 Azure Backup Azure Backup .

VM1 .

VM1 .

? .

: 1 .

You can perform a file recovery of VM1 to:

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

You can restore VM1 to:

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

Answer:

You can perform a file recovery of VM1 to:

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

You can restore VM1 to:

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

Explanation:

Box 1 : VM1 and VM2 only

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

Therefore -

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

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

Box 2 : VM1 or A new Azure virtual machine only

When restoring a VM, you can't use the replace existing VM option for encrypted VMs. This option is only supported for unencrypted managed disks. And also You can restore files from a VM to the same server operating system, or to the compatible client operating system only. Hence " VM1 or A new Azure virtual machine only " is correct answer.

Answer Area



You can perform a file recovery of VM1 to:

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

You can restore VM1 to:

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

References:

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-arm-restore-vm>

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-restore-files-from-vm#system-requirements>

NEW QUESTION: 72

VM1 is affected by planned maintenance. VM1 ARM1.json is an ARM template that is used to create VM1 in the resource group. VM1 is affected by planned maintenance, and the goal is to move VM1 to a different host immediately. According to Microsoft Azure Administrator documentation, which of the following are supported options to move VM1 to a different host immediately?

- A.
- B.

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: 73

Scenario: A company has an Azure subscription. The company wants to create a Logic App in the subscription. The Logic App will be used to trigger an Azure Function when a new resource is created in the subscription. The Logic App will also be used to trigger an Azure Function when a new resource is deleted in the subscription.

The company has an Azure AD tenant. The company wants to create a role assignment for the Logic App. The role assignment will grant the Logic App Contributor permissions on the subscription.

Adatum has an Azure AD tenant. Adatum has a subscription named Subscription1. Adatum has a user named Developers. Adatum has a group named DevOps. Adatum has a role assignment for the Logic App Contributor role on the subscription. The role assignment is assigned to the DevOps group.

DevOps has a Logic App named LogicApp1. LogicApp1 is used to trigger an Azure Function when a new resource is created in the subscription.

LogicApp1 is used to trigger an Azure Function when a new resource is deleted in the subscription.

LogicApp1 is used to trigger an Azure Function when a new resource is created in the subscription?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

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

[Azure Logic Apps permissions and access control]

NEW QUESTION: 74

Scenario: A company has an Azure subscription. The company wants to monitor network connectivity and latency between Azure virtual machines and on-premises resources using Azure Network Watcher - Connection Monitor (v2).

The company has an Azure subscription. The company has a virtual machine named VM1 in the subscription. The company has a domain controller named DC1 in the datacenter. The company has an ExpressRoute connection between the subscription and the datacenter.

VM1 and DC1 are connected via ExpressRoute. The company wants to monitor network connectivity and latency between VM1 and DC1 using Connection Monitor (v2).

DC1 is connected to the subscription via ExpressRoute?

A. Log Analytics agents

B. Azure Network Watcher Agent on VM1 and DC1

C. Azure Monitor agents on VM1 and DC1

D. Azure Arc agents on VM1 and DC1

Answer: D (LEAVE A REPLY)

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

Scenario Breakdown

VM1: Azure virtual machine (in your subscription)

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

Connectivity: Via ExpressRoute

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

Understanding Azure Connection Monitor (v2)

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

Number of virtual networks: ▼

1
2
3

Number of subnets: ▼

1
2
3

Answer:

Number of virtual networks: ▼

1
2
3

Number of subnets: ▼

1
2
3

Explanation:

Number of virtual networks: ▼

1
2
3

Number of subnets: ▼

1
2
3

In this scenario, App1 consists of three distinct tiers - web front end, processing middle tier, and SQL database - each containing five virtual machines. The technical requirement specifies that the company must minimize the number of open ports between the App1 tiers, move all tiers of App1 to Azure, and ensure that all VMs are protected by backups.

According to Microsoft Azure architecture best practices for multi-tier applications (from Azure Architecture Center and the Azure Administrator curriculum), the optimal design involves:

- * Deploying all tiers of App1 into a single virtual network (VNet).
- * This allows all components of the application to communicate securely using private IP addresses.
- * Keeping all tiers within a single VNet simplifies management, security, and monitoring while supporting Network Security Groups (NSGs) for inter-tier traffic control.

* Microsoft Documentation Extract:

"Use a single virtual network to host multi-tier applications. Divide the virtual network into multiple subnets, each representing a tier, and use network security groups (NSGs) to control traffic flow between tiers." (Source: Microsoft Learn - Design and implement virtual networks in Azure)

- * Creating separate subnets for each application tier (3 total).

Which IP address can be assigned to an Azure Firewall Premium SKU?

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

FW1 is an Azure Firewall Premium SKU. Which IP address can be assigned to it?

Which IP address can be assigned to it?

- A. IP2
- B. IP1 or IP2
- C. IP1, IP2, IP5
- D. IP1, IP2, IP4, IP5

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: 78

Which Azure service can be used to create a deny assignment?

storage1 is a storage account. Which Azure service can be used to create a deny assignment?

storage1 is a storage account. Which Azure service can be used to create a deny assignment?

Which Azure service can be used to create a deny assignment?

- A. Azure Resource Manager (ARM)
- B. Azure
- C. Azure Policy
- D. Azure Security Center

Answer: (SHOW ANSWER)

A deny assignment prevents users from modifying or deleting resources, even if they have RBAC permissions. Microsoft explicitly documents that deny assignments are automatically created and managed by Azure deployment stacks.

Azure Policy enforces compliance but does not create deny assignments. ARM templates deploy resources but cannot create deny assignments directly. Landing zones are architectural patterns, not enforcement mechanisms.

Microsoft documentation states:

"Deployment stacks create and manage deny assignments to prevent unauthorized changes to managed resources."

NEW QUESTION: 79

www.contoso.com is a CNAME record. Which Azure service can be used to create a deny assignment?

Which Azure service can be used to create a deny assignment?

- A. Azure Resource Manager (ARM)
- B. www.contoso.com asuid.contoso.com
- C. Azure Policy
- D. www.contoso.com asuid.contoso.com

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: 80

VM1 VM2 Azure App1

App Azure VM1 VM2 App1

?

A.

B.

C.

D.

Answer: (SHOW ANSWER)

An Availability Set in Azure is a logical grouping of two or more virtual machines that helps ensure your application remains available during planned maintenance events and hardware failures within Azure datacenters.

According to the Microsoft Azure Administrator (AZ-104) study guide and Azure official documentation, Availability Sets protect your VMs from both:

Hardware failures (through fault domains)

Planned maintenance events (through update domains)

A Fault Domain (FD) represents a group of virtual machines that share a common power source and physical network switch. By default, Azure assigns up to three fault domains per region to distribute resources across different racks or hardware clusters.

An Update Domain (UD) represents a group of virtual machines that can be rebooted together during planned maintenance by Microsoft. Azure never updates or reboots two update domains at the same time, ensuring high availability during maintenance operations.

When you have two VMs (VM1 and VM2) running the same application (App1), you should place them in an Availability Set with at least two update domains. This ensures that when Azure performs planned maintenance (updates to host OS or hypervisor), only one update domain is taken offline, while the other remains operational - maintaining service continuity for App1.

By default, an Availability Set provides 5 update domains and 3 fault domains, but the key requirement here is at least two update domains to ensure that maintenance on one host does not impact the entire application.

Therefore, to keep App1 available during planned maintenance of the hardware hosting VM1 and VM2, the Availability Set must include two update domains so that each VM resides in a separate update domain.

NEW QUESTION: 81

Which of the following is a valid Recovery Services vault name? (Select the correct answer.)

Recovery Services vault names must be unique within a subscription.

Which of the following is a valid Recovery Services vault name?

- A. Recovery Services - My Vault
- B. Recovery Services - My Vault - 1234
- C. Recovery Services - My Vault - 1234 - 5678
- D. Recovery Services - My Vault - 1234 - 5678 - 9010

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 82

You are configuring an Azure tenant for Microsoft Entra ID. The tenant has the following structure:

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

The tenant has the following subscriptions:

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

The tenant has the following users and roles:

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

Which of the following statements are true? (Select all that apply.)

Answer Area

Statements	Yes	No
User1 can resize VM1.	<input type="radio"/>	<input type="radio"/>
User2 can create a new storage account in RG1.	<input type="radio"/>	<input type="radio"/>
User3 can assign User1 the Owner role for RG3.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
User1 can resize VM1.	<input checked="" type="radio"/>	<input type="radio"/>
User2 can create a new storage account in RG1.	<input type="radio"/>	<input checked="" type="radio"/>
User3 can assign User1 the Owner role for RG3.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
User1 can resize VM1.	<input checked="" type="radio"/>	<input type="radio"/>
User2 can create a new storage account in RG1.	<input type="radio"/>	<input checked="" type="radio"/>
User3 can assign User1 the Owner role for RG3.	<input checked="" type="radio"/>	<input type="radio"/>

This question is about how RBAC scope inheritance and built-in role permissions work across management groups, subscriptions, and resource groups. In Azure, role assignments made at a management group scope apply to all child subscriptions and their resources through inheritance. Here, Sub3 is a child of MG2, and RG3 (which contains VM1) is in Sub3. Since User1 is assigned the Contributor role at scope MG2, User1 inherits Contributor permissions on Sub3 and RG3. The Contributor role grants management permissions to create, update, and delete Azure resources (but not grant access). Resizing a VM is a management operation on the VM resource, so User1 can resize VM1 (Yes).

User2 is assigned Storage Account Contributor at the scope of the existing storage1 resource only. That role allows management actions for that specific storage account (such as configuration) but does not grant permissions to create new storage accounts in the resource group unless assigned at RG/subscription scope.

Therefore, User2 cannot create a new storage account in RG1 (No).

User3 has User Access Administrator at the Tenant Root Group. This role is specifically for managing access (role assignments) across the assigned scope and its children, allowing User3 to assign roles like Owner at RG3. Thus, User3 can assign User1 the Owner role for RG3 (Yes).

NEW QUESTION: 83

VMSS1 is a Virtual Machine Scale Set in Azure. Which of the following are valid placement groups for VMSS1?

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

VMSS1 is a Virtual Machine Scale Set in Azure. Which of the following are valid placement groups for VMSS1?

VMSS1 is a Virtual Machine Scale Set in Azure. Which of the following are valid placement groups for VMSS1?

- A. Proximity1
- B. Proximity1, Proximity2, Proximity3
- C. Proximity1 and Proximity3
- D. Proximity2

Answer: A (LEAVE A REPLY)

Placement Groups is a capability to achieve co-location of your Azure Infrastructure as a Service (IaaS) resources and low network latency among them, for improved application performance.

Azure proximity placement groups represent a new logical grouping capability for your Azure Virtual Machines, which in turn is used as a deployment constraint when selecting where to place your virtual machines. In fact, when you assign your virtual machines to a proximity placement group, the virtual machines are placed in the same data center, resulting in lower and deterministic latency for your applications.

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

Reference:

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

NEW QUESTION: 84

Azure storage1 is a Storage Account in Azure. Which of the following are valid placement groups for storage1?

90□ □□ □□□ □□□ □□□□ □□ □□ □□ □□□□ □□□□ storage1□ □□ □□ □□ □□ □□□ □□□□ □□□.
□□□ □□□ □□□□ □□□? □□□□ □□ □□□□ □□□ □□□ □□□□□.
□□: □□ 1□□ 1□□□□.

Answer Area

```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToArchive": {
              "enableAutoTierToHotFromCool": {
                "tierToArchive": {
                  "tierToCool": {
                    "daysAfterModificationGreaterThan": 90
                  }
                }
              }
            }
          }
        }
      }
    }
  ]
}
```



Answer:



```
{
  "rules": [
    {
      "enabled": true,
      "name": "rule1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToArchive": {
              "enableAutoTierToHotFromCool": {
                "tierToArchive": {
                  "tierToCool": {
                    "daysAfterModificationGreaterThan": 90
                  }
                }
              }
            }
          }
        }
      }
    }
  ]
}
```

Explanation:

According to the Microsoft Azure Administrator documentation, the Custom Script Extension is used to download and execute scripts on Azure virtual machines after deployment. It is ideal for configuration management, application installation, and post-deployment software setup.

The extension can execute PowerShell or Bash scripts stored in Azure Storage or GitHub repositories. When used with an ARM template, you define the extension in the template's "resources" section under "type": "Microsoft.Compute/virtualMachines/extensions" or "Microsoft.Compute/virtualMachineScaleSets/extensions".

Microsoft's official documentation notes:

"The Custom Script Extension for Windows and Linux downloads and executes scripts on Azure virtual machines. This extension is useful for post-deployment configuration, software installation, or any other configuration or management tasks." In this case, using Azure Custom Script Extension allows automatic installation of NGINX on all VM instances as soon as they are created in the scale set - satisfying the requirement that NGINX is available immediately after deployment.

Other options do not fit:

* Application Insights is used for monitoring, not configuration.

* Publish-AzVMDscConfiguration and New-AzConfigurationAssignment are used for Desired State Configuration (DSC), not lightweight application installation like NGINX.

Hence, the only verified and cost-efficient solution is Azure Custom Script Extension.

NEW QUESTION: 86

RG1 is a resource group in Subscription1. Azure LB1 is a load balancer.

RG1 contains LB1 and LB2. LB2 is a load balancer.

Admin1 is a user who has the Network Contributor role on LB1 and LB2. Admin1 is also a user who has the Contributor role on RG1.

Admin1 wants to add a health probe to LB2. What is the minimum role assignment that Admin1 needs to have on RG1 to perform this task?

Options:
A. Contributor
B. Network Contributor
C. Owner
D. Network Contributor on LB2



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: 87

VM1 is a Windows VM in a resource group. You need to ensure that VM1 can access the Azure Key Vault service.

VM1 is a Linux VM in a resource group. You need to ensure that VM1 can access the Azure Key Vault service.

* VM1 is a Windows VM in a resource group. You need to ensure that VM1 can access the Azure Key Vault service.

* VM1 is a Linux VM in a resource group. You need to ensure that VM1 can access the Azure Key Vault service.

* VM1 is a Windows VM in a resource group. You need to ensure that VM1 can access the Azure Key Vault service.

Which role assignment should you use to ensure that VM1 can access the Azure Key Vault service?

A. Network Contributor on the resource group

B. Contributor on the resource group

C. Azure Key Vault Contributor on the resource group

D. Contributor on the Azure Key Vault service

Answer: (SHOW ANSWER)

Azure Disk Encryption is a service that helps you encrypt your Windows and Linux IaaS virtual machine disks¹. It uses BitLocker for Windows and DM-Crypt for Linux to provide volume encryption for the OS and data disks². Azure Disk Encryption requires that you use a key encryption key in Azure Key Vault to encrypt the volume encryption key, which is then stored on the disk. You can use either a service-managed key or a customer-managed key in Azure Key Vault³. Azure Disk Encryption also supports encrypting virtual machine disks that are downloaded from Azure⁴.

NEW QUESTION: 88

VM4 is configured with an Azure Event Grid subscription. You want to create a workflow to send an email message when the settings of VM4 are modified.

- A. Azure Logic App
- B. Azure Event Grid
- C. Azure Logic App
- D. Azure Event Grid

Answer: B (LEAVE A REPLY)

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

You can start an automated logic app workflow when specific events happen in Azure resources or third-party resources. These resources can publish those events to an Azure event grid. In turn, the event grid pushes those events to subscribers that have queues, webhooks, or event hubs as endpoints. As a subscriber, your logic app can wait for those events from the event grid before running automated workflows to perform tasks

- without you writing any code.

References:

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

NEW QUESTION: 89

You have three Azure Storage accounts named storage1, storage2, and storage3. You want to configure lifecycle management for storage1, storage2, and storage3.

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

You want to configure lifecycle management for storage1, storage2, and storage3. You want to configure lifecycle management for storage1, storage2, and storage3. You want to configure lifecycle management for storage1, storage2, and storage3.

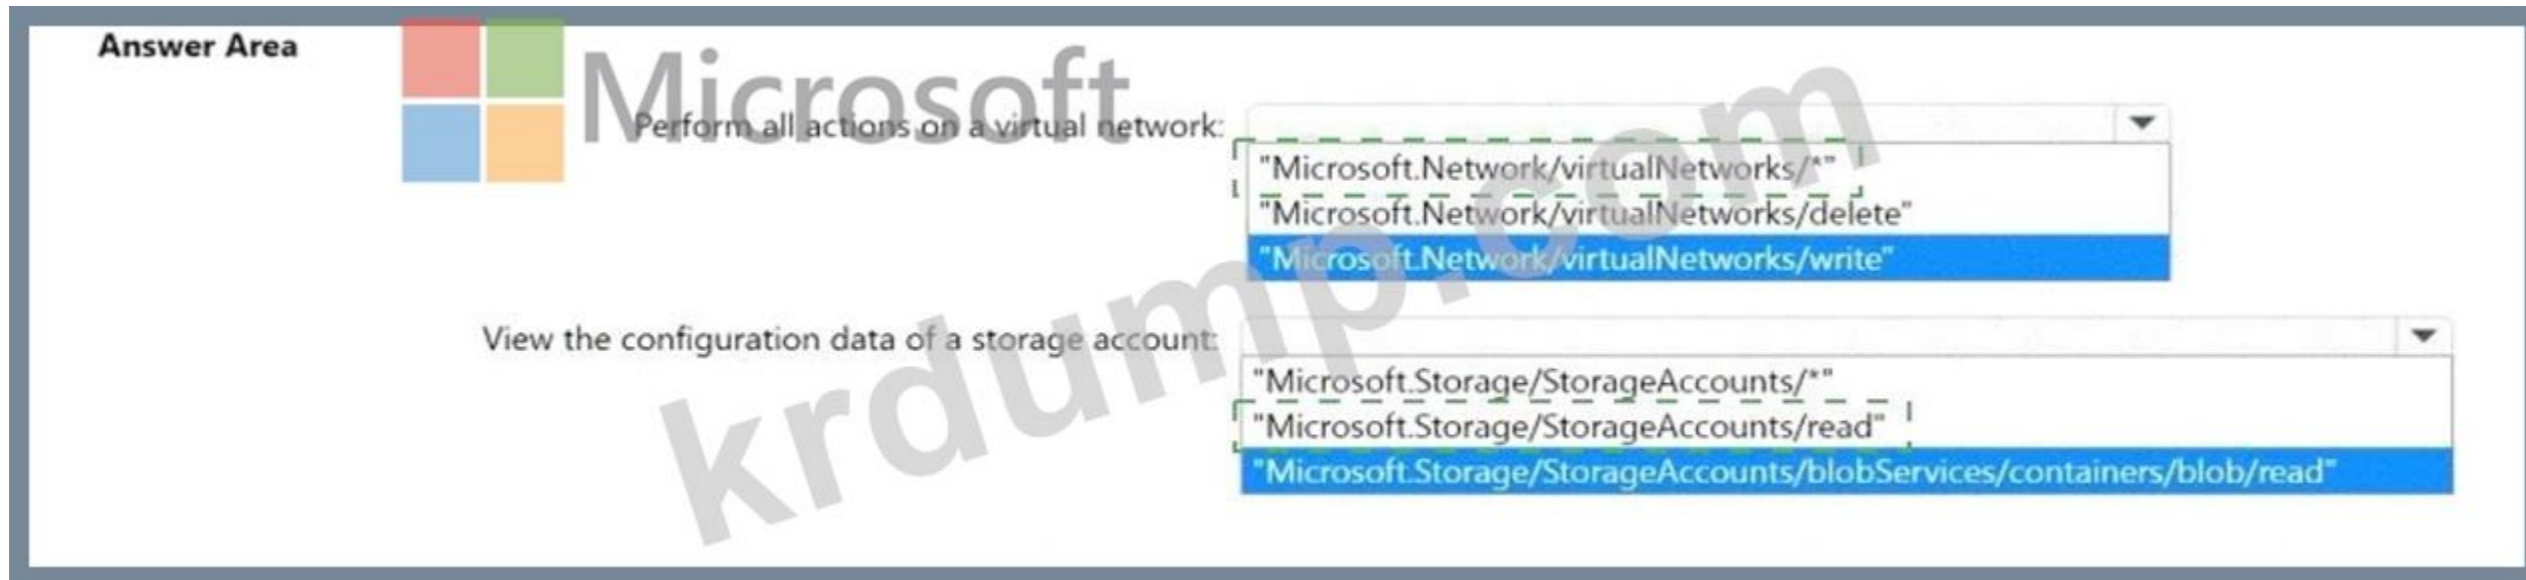
Lifecycle management:

- storage1 only
- storage2 only
- storage1 and storage3 only
- storage2 and storage3 only
- storage1, storage2, and storage3

The Archive access tier:

- storage1 only
- storage2 only
- storage1 and storage3 only
- storage2 and storage3 only
- storage1, storage2, and storage3

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: 91

Subscription1 is an Azure subscription.

Subscription1 contains the following resources:

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

Subscription1 contains an Azure Resource Manager resource group named RG1. RG1 contains the following resources:

VM1 (Virtual Machine)

Storage2 (Azure Storage account)

A. RG1

B. VM1

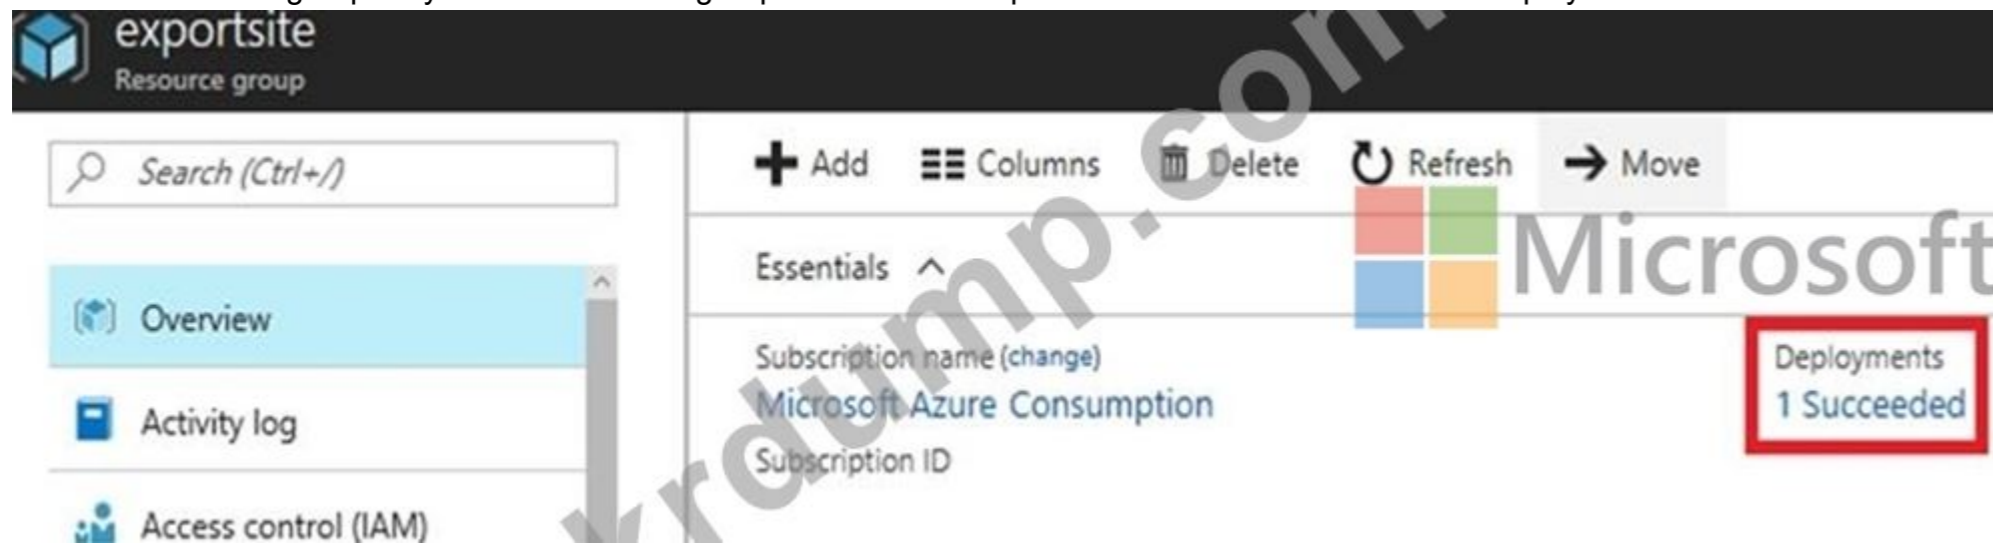
C. Storage1

D. Storage2

Answer: A (LEAVE A REPLY)


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.

To inspect or capture network traffic between two Azure virtual machines (VM1 and VM2) within the same virtual network, Azure provides specialized network diagnostic tools - Azure Network Watcher, Connection Monitor, and Packet Capture - rather than Windows Performance Monitor.

The Performance Monitor (PerfMon) tool on Windows is used to collect metrics such as CPU, memory, and disk I/O performance counters, not network packet data or traffic analysis between VMs.

While you could use PerfMon to track bandwidth usage or network throughput locally, it cannot capture or inspect packet-level communication or determine the content or flow between two Azure VMs.

According to the Microsoft Azure Administrator documentation (AZ-104 study guide), the correct method to analyze and inspect network traffic between Azure virtual machines includes:

Enabling Azure Network Watcher in the target region.

Using Packet Capture under Network Watcher to capture inbound and outbound packets.

Optionally using Connection Monitor to monitor connectivity metrics between endpoints.

Packet Capture allows you to define capture filters (e.g., source/destination IP, ports, protocols) and run it for a specific duration, such as three hours, as required in this scenario. The capture file is stored in Azure Storage or locally for later analysis with tools like Wireshark.

Therefore, creating a Data Collector Set (DCS) in Performance Monitor does not meet the requirement to inspect network traffic between VM1 and VM2.

NEW QUESTION: 93

□□□□ □□□ Azure □□□ □□□□□.

□□ □□□ □□ Deploy json□□□ □□□ □□□□□.

```

1 {
2   "$schema": "https://schema.management.azure.com/schemas/2015-01-01/deploymentTemplate.json#",
3   "contentVersion": "1.0.0.0",
4   "parameters": {},
5   "variables": {},
6   "resources": [
7     {
8       "type": "Microsoft.Resources/resourceGroups",
9       "apiVersion": "2018-05-01",
10      "location": "eastus",
11      "name": "[concat('RG', copyIndex())]",
12      "copy": {
13        "name": "copy",
14        "count": 3
15      }
16    },
17    {
18      "type": "Microsoft.Resources/deployments",
19      "apiVersion": "2021-04-01",
20      "name": "lockDeployment",
21      "resourceGroup": "RG1",
22      "dependsOn": [[resourceId('Microsoft.Resources/resourceGroups/', 'RG1')]],
23      "properties": {
24        "mode": "Incremental",
25        "template": {
26          "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
27          "contentVersion": "1.0.0.0",
28          "parameters": {},
29          "variables": {},
30          "resources": [
31            {
32              "type": "Microsoft.Authorization/locks",
33              "apiVersion": "2016-09-01",
34              "name": "rglock",
35              "properties": {
36                "level": "CanNotDelete"
37              }
38            }
39          ]
40        }
41      }
42    },
43    {
44      "type": "Microsoft.Resources/deployments",
45      "apiVersion": "2021-04-01",
46      "name": "lockDeployment",
47      "resourceGroup": "RG2",
48      "dependsOn": [[resourceId('Microsoft.Resources/resourceGroups/', 'RG2')]],
49      "properties": {
50        "mode": "Incremental",
51        "contentVersion": "1.0.0.0",
52        "parameters": {},
53        "variables": {},
54        "resources": [
55          {
56            "type": "Microsoft.Authorization/locks",
57            "apiVersion": "2016-09-01",
58            "name": "rgLock",
59            "properties": {
60              "level": "ReadOnly"
61            }
62          }
63        ]
64      }
65    }
66  ]
67 }
68 }
69 }
70 "outputs": {}
71 }

```

□□□ □□□□ □□ cmdlet□ □□□□□.

New-AzDeployment -□□ westus -□□□ □□ "deploy.json"

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

Answer Area

Statements	Yes	No
You can deploy a virtual machine to RG1.	<input type="radio"/>	<input type="radio"/>
You can deploy a virtual machine to RG2.	<input type="radio"/>	<input type="radio"/>
You can manually create a resource group named RG3.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
You can deploy a virtual machine to RG1.	<input type="radio"/>	<input type="radio"/>
You can deploy a virtual machine to RG2.	<input type="radio"/>	<input type="radio"/>
You can manually create a resource group named RG3.	<input type="radio"/>	<input type="radio"/>

Explanation:

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

You can deploy a virtual machine to RG1. = No

You can deploy a virtual machine to RG2. = No

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

Let me explain why:

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

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

If you enter RG1 as the value for the resourceGroupName parameter and VM1 as the value for the vmName parameter, then the cmdlet will create a resource group named RG1 and a virtual machine named VM1 in the westus location. Therefore, you can deploy a virtual machine to RG1.

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

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

This cmdlet takes two parameters: Name and Location. For example, you can run the following cmdlet to create a resource group named RG3 in westus:

New-AzResourceGroup -Name RG3 -Location westus

NEW QUESTION: 94

□□ □□ □□□ Azure □□□ □□ □□□□□.

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

How many App Service plans are required to host all four web applications?

- 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: 95

Which of the following is a valid configuration for peering two VNet1 and VNet2 in Azure?

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

Which of the following is a valid configuration for peering two VNet1 and VNet2 in Azure?

Azure Bastion is required for peering VNet1 and VNet2.

* VNet1 and VNet2 must be in the same region.

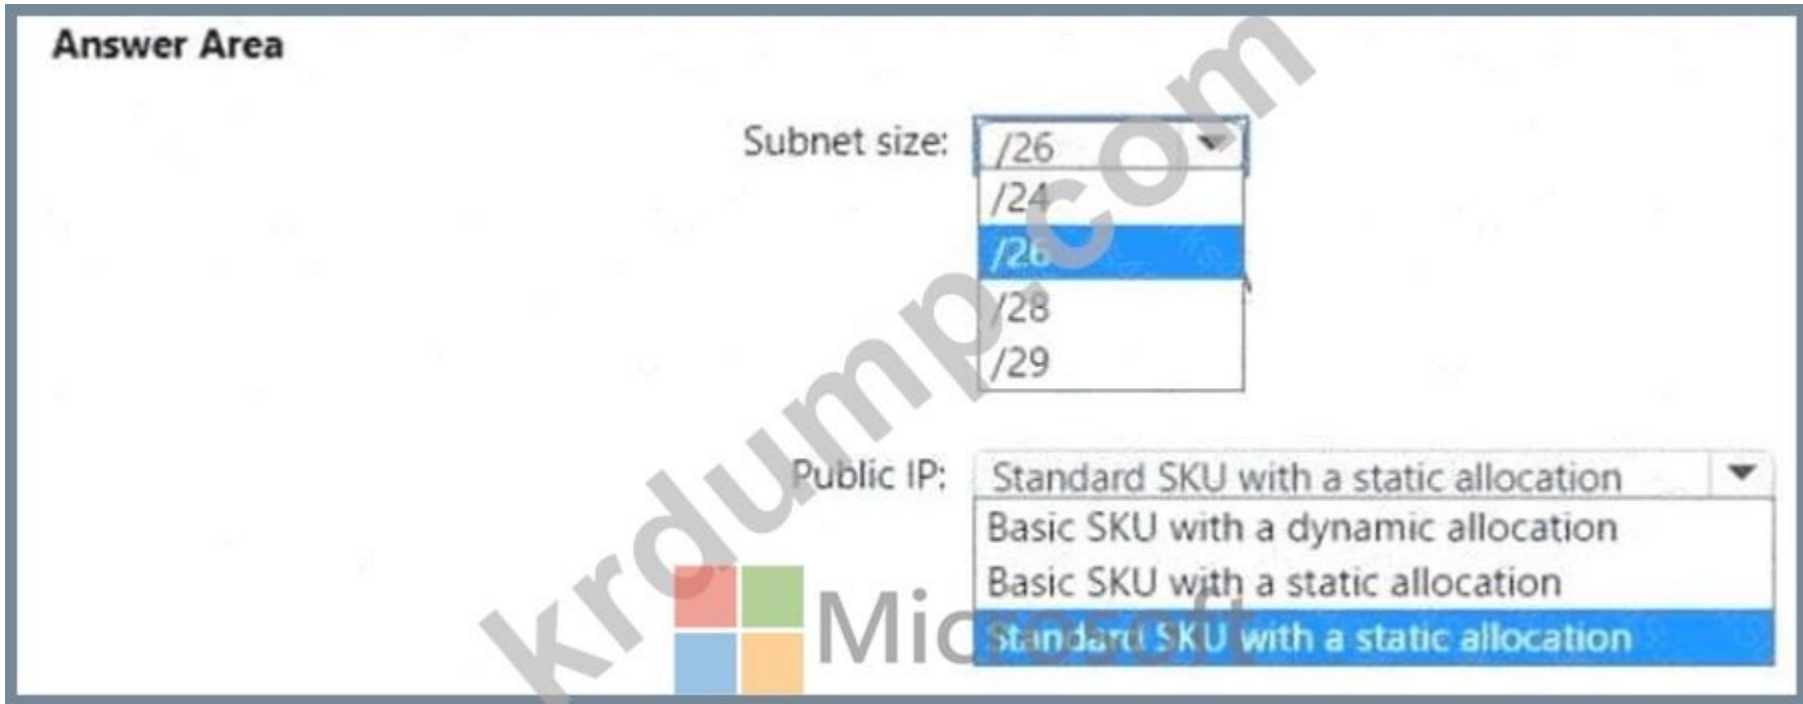
* VNet1 and VNet2 must be in the same address space.

* VNet1 and VNet2 must be in the same virtual address space.

* Azure Bastion is required for peering VNet1 and VNet2.

Azure Bastion is required for peering VNet1 and VNet2? VNet1 and VNet2 must be in the same region.

Answer: VNet1 and VNet2 must be in the same region.



Answer:



Explanation:



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

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

1## Support for host scaling:

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

2## Support uploading and downloading files:

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

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

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

4## Minimize the number of addresses on the Azure Bastion subnet:

Azure Bastion requires a dedicated subnet named AzureBastionSubnet.

* The minimum supported subnet size is /26 for the Standard SKU (as it supports scaling and multiple instances).

* The Basic SKU can use /27, but since we are using Standard SKU (for scaling and file transfer), the minimum possible subnet size is /26. This meets the requirement to minimize address space usage while supporting scaling.

5## Public IP requirements:

* The Standard SKU Bastion requires a Public IP address of SKU = Standard with Static allocation.

* Basic SKU Bastion can work with Basic Public IPs, but not Standard SKU Bastion. Hence, we must use a Standard SKU Public IP with Static allocation.

Final Verified Configuration (per Microsoft Azure Administrator Documentation):

* Subnet size: /26

* Public IP: Standard SKU with a static allocation

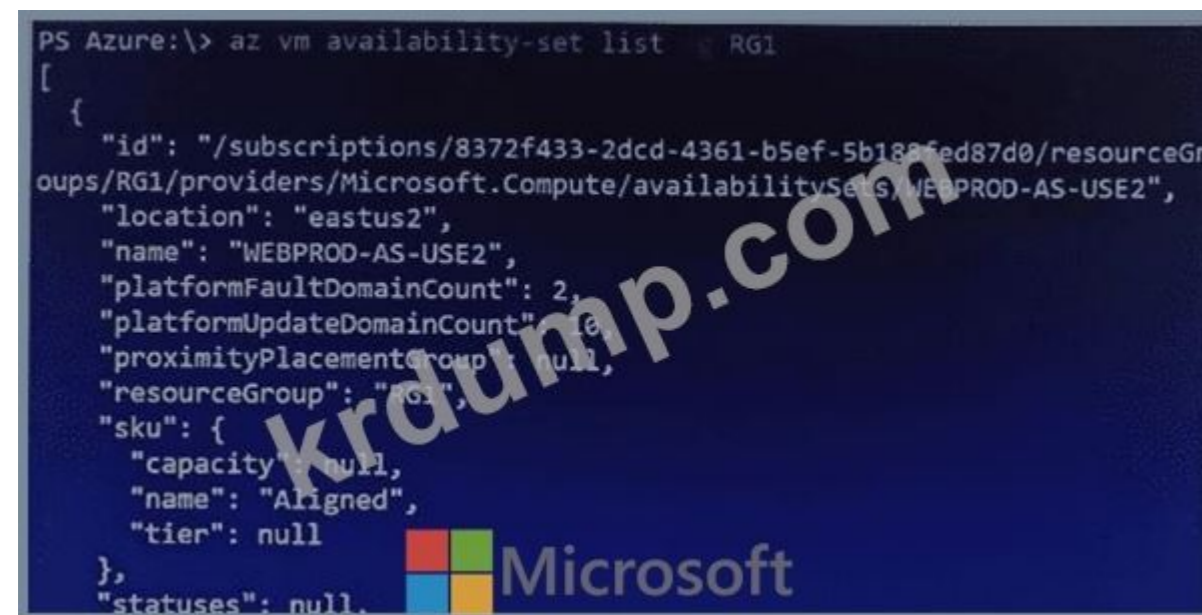
Rationale Summary:

This configuration supports scaling, file transfer, cross-VNet connectivity, and minimal address consumption, satisfying all requirements as per official Azure documentation on Azure Bastion Standard SKU and Bastion network design guidelines.

NEW QUESTION: 96

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

```
PS Azure:\> az vm availability-set list --rg RG1
[
  {
    "id": "/subscriptions/8372f433-2dcd-4361-b5ef-5b188fed87d0/resourceGroups/RG1/providers/Microsoft.Compute/availabilitySets/WEBPROD-AS-USE2",
    "location": "eastus2",
    "name": "WEBPROD-AS-USE2",
    "platformFaultDomainCount": 2,
    "platformUpdateDomainCount": 16,
    "proximityPlacementGroup": null,
    "resourceGroup": "RG1",
    "sku": {
      "capacity": null,
      "name": "Aligned",
      "tier": null
    },
    "statuses": null
  }
]
```



```
"tags": {},
"type": "Microsoft.Compute/availabilitySets",
"virtualMachines": [
]
}
Azure/ Microsoft
```

WEBPROD-AS-USE2 14 VMs are distributed across 10 update domains. When Microsoft performs planned maintenance in East US 2, the maximum number of unavailable virtual machines will be [answer choice].
If the server rack in the Azure datacenter that hosts WEBPROD-AS-USE2 experiences a power failure, the maximum number of unavailable virtual machines will be [answer choice].

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

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

Answer:

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

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

Explanation:

Box 1: 2

There are 10 update domains. The 14 VMs are shared across the 10 update domains so four update domains will have two VMs and six update domains will have one VM. Only one update domain is rebooted at a time.

Therefore, a maximum of two VMs will be offline.

Box 2: 7

There are 2 fault domains. The 14 VMs are shared across the 2 fault domains, so 7 VMs in each fault domain. A rack failure will affect one fault domain so 7 VMs will be offline.

Answer Area



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

	▼
2	
7	
10	
14	

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

	▼
2	
7	
10	
14	

Reference:

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

NEW QUESTION: 97

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

□□ □□ □□□ □□□□ □□□□?

A. □□ □□

B. Azure □□ □□

C. □□ □□

D. □□ □□□ □□□

Answer: D (LEAVE A REPLY)

A Recovery Services vault is a logical container that stores the backup data for each protected resource, such as Azure VMs. When the backup job for a protected resource runs, it creates a recovery point inside the Recovery Services vault.

Scenario:

There are three application tiers, each with five virtual machines.

Move all the virtual machines for App1 to Azure.

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

References: <https://docs.microsoft.com/en-us/azure/backup/quick-backup-vm-portal>

NEW QUESTION: 98

Template1□ Template2□□ □□□ Azure Resource Manager(ARM) □□□□ □ □ □□□□.

□□ Template1 □□□ □□□□ □□□□ Template1□ □□ □□ □□ □□□□ Template2□□ □□□□ □□□□ □□□□ □□□□□□.

Template1□ Template2□ □□ □□□□□□ □□ □□□□□□ □□, □□ □□□ □□□ □□□ □□ □□ □□□ □□□□ □□□□ □□□□ □□□□ □□□□ □□□□.

□□□ □□ □□□?

- A. Template1□□ □□ □□ □□ □□□□□ provisionAfterExtensions □□□ □□□□□.
- B. Template1□□ □□ □□ □□ □□□□□ dependsOn □□□ □□□□□.
- C. Template2□ □□□□ Template1□ □□□□□.
- D. Template1□ Template2□ □□□ □□□□□ □□□□ □□ □□□□ □□□□□.

Answer: D (LEAVE A REPLY)

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

The correct approach is to create a parent ARM template that references Template1 and Template2 as linked templates . Microsoft's ARM guidance states that complex solutions can be broken into related templates and deployed together through a main template; linked templates remain separate files referenced from the main template, which preserves modularity and reusability.

This also solves the cross-template dependency issue. Microsoft states that dependencies can be set between nested or linked templates, and if resources in one linked template must exist before resources in another, the second linked deployment can depend on the first. A simple dependsOn on the VM extension inside Template1 is insufficient because ARM dependencies are evaluated for resources deployed in the same template context; the documented dependency behavior is that Resource Manager deploys resources in dependency order and otherwise deploys them in parallel.

Adding Template2's resources into Template1 would violate the requirement that both templates remain separate and reusable. This aligns with AZ-104 Deploy and manage Azure compute resources , specifically Automate deployment of resources by using Azure Resource Manager templates or Bicep files .

NEW QUESTION: 99

Windows Server 2019□ □□□□ □□ □□ □□ □□□ Azure □□ □□□ □□□□.

Name	Private IP address	Public IP address	Virtual network name	DNS suffix configured in Windows Server
VM1	10.1.0.4	52.186.85.63	VNET1	Adatum.com
VM2	10.1.0.5	13.92.168.13	VNET1	Contoso.com

adatum.com□□□ □□□ □□□□ Azure DNS □□□ □□□□. VNET1□□ □□ □□□ □□□□□ adatum.com □□□ □□□□□.

□ □□ □□□ adatum.com □□□□ □□ A □□□□ □□□□□? □□□□ □□ □□□□ □□□ □□□ □□□□□.

□□: □□ □□□ 1□□□□□.



Answer:

Answer Area



Explanation:



NEW QUESTION: 100

Contoso.com Azure Active Directory. Contoso.com Windows 10.

Name	Role
User1	Cloud device administrator
User2	User administrator

Contoso.com Windows 10.

Name	Join type
Device1	Azure AD registered
Device2	Azure AD joined

Contoso.com Windows 10.

Name	Join type	Owner
Group1	Assigned	User1
Group2	Dynamic Device	User2

Group1, Group2, Device1, Device2. Group1, Group2, Device1, Device2.

Statements	Yes	No
User1 can add Device2 to Group1	<input type="radio"/>	<input type="radio"/>
User2 can add Device1 to Group1	<input type="radio"/>	<input type="radio"/>
User2 can add Device2 to Group2	<input type="radio"/>	<input type="radio"/>

Answer:

Statements Yes No

User1 can add Device2 to Group1	<input checked="" type="radio"/>	<input type="radio"/>
User2 can add Device1 to Group1	<input type="radio"/>	<input checked="" type="radio"/>
User2 can add Device2 to Group2	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Statements	Yes	No
User1 can add Device2 to Group1	<input type="radio"/>	<input type="radio"/>
User2 can add Device1 to Group1	<input type="radio"/>	<input checked="" type="radio"/>
User2 can add Device2 to Group2	<input checked="" type="radio"/>	<input type="radio"/>

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: 101

RSV1 Recovery Services 5 VMs. RSV1 5 VMs. RSV1 14 VMs. RSV1 14 VMs.

RSV1 VM1 8 VMs. VM1 8 VMs. VM1 8 VMs. VM1 8 VMs.

VM1 8 VMs. VM1 8 VMs. VM1 8 VMs. VM1 8 VMs.

VM1 8 VMs?

A. VM1 8 VMs.

B. VM1 8 VMs. VM1 8 VMs.

C. VM1 .

D. VM1 .

Answer: (SHOW ANSWER)

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

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

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

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

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

Here's how it works:

The backup is used to create a new VM in the same or different resource group.

Once the new VM is confirmed operational, you can redirect traffic or replace the original VM during a controlled maintenance window.

The "Replace existing" option, on the other hand, overwrites the current VM - leading to downtime and risk if validation fails.

Therefore, to restore VM1 to its state from eight days ago while minimizing downtime, you must first restore using the "Create new restore configuration" option, validate the restore, and then switch over seamlessly.

NEW QUESTION: 102

Azure .

Name	Operating system
Image1	Windows Server
Image2	Linux

.

* Azure

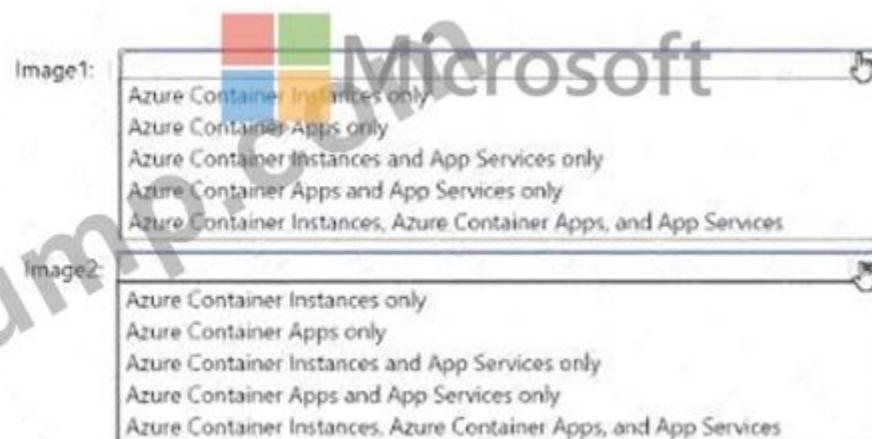
* Azure

* Azure

? .

: 1 .

Answer Area



Answer:

Answer Area



Explanation:

Image 1: Azure Container Apps only. image 2: Azure Container Instances, Azure Container Apps, and App Services.

The images you have in your Azure subscription are different types of container images that can run on different Azure services. A container image is a package of software that includes everything needed to run an application, such as code, libraries, dependencies, and configuration files. Container images are portable and consistent across different environments, such as development, testing, and production.

Azure Container Instances is a service that allows you to run containers directly on the Azure cloud, without having to manage any infrastructure or orchestrators. You can use Azure Container Instances to run any container image that is compatible with the Docker image format and follows the Open Container Initiative (OCI) specification. You can also run Windows or Linux containers on Azure Container Instances.

Azure Container Apps is a service that allows you to build and deploy cloud-native applications and microservices using serverless containers. You can use Azure Container Apps to run any container image that is compatible with the Docker image format and follows the Open Container Initiative (OCI) specification.

You can also run Windows or Linux containers on Azure Container Apps.

Azure App Service is a service that allows you to build and host web applications, mobile backends, and RESTful APIs using various languages and frameworks. You can use Azure App Service to run custom container images that are compatible with the Docker image format and follow the App Service Docker image contract. You can also run Windows or Linux containers on Azure App Service.

NEW QUESTION: 103

storage1 Azure .

Blob (SAS) .

? .

: 1 .

Answer Area

Allowed services ⓘ

Blob File Queue Table

Allowed resource types ⓘ

Service Container Object

Allowed permissions ⓘ

Read Write Delete List Add Create Update Process Immutable storage Permanent delete

Blob versioning permissions ⓘ

Enables deletion of versions

Allowed blob index permissions ⓘ

Read/Write Filter

Start and expiry date/time ⓘ

Start

End

(UTC) Coordinated Universal Time

Allowed IP addresses ⓘ

Allowed protocols ⓘ

HTTPS only HTTPS and HTTP



Krdump.com

Answer:

* 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: 105

VM1 is an Azure VM with the following configuration:

OS: Windows Server 2016
Storage: 100 GB
Network: 1 NIC

VM1 is in a state that prevents backup. What is the cause?

A. VM1 is in a state that prevents backup.

B. VM1 is in a state that prevents backup.

C. VM1 is in a state that prevents backup because the latest VM Agent (WaAppAgent.exe) is not installed.

D. VM1 is in a state that prevents backup.

Answer: C (LEAVE A REPLY)

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

Azure Backup Pre-Check validates whether the VM configuration is suitable for successful backup before or during backup processing. In Microsoft's description of Azure VM backup pre-checks, a Warning state indicates "one or more issues in VM's configuration that might lead to backup failures," and Microsoft specifically gives the example that not having the latest VM Agent installed can cause intermittent backup failures and falls into this Warning category.

Therefore, the correct answer is C. The Azure VM Agent, shown on Windows VMs as WaAppAgent.exe, enables communication between the Azure platform, VM extensions, and backup components.

Azure Backup relies on this integration to install and manage the VM snapshot extension and coordinate backup operations.

Microsoft troubleshooting guidance also lists ensuring that the VM Agent is the latest version as a core backup troubleshooting requirement.

The other options do not best match a Warning pre-check status. A vault availability issue is not a VM configuration warning. An unmanaged disk is not the expected pre-check warning cause here. A stopped VM state alone is not the cited reason for this warning. Study Guide reference: AZ-104 "Monitor and back up Azure resources" > "Implement backup and recovery" > Azure Backup / Recovery Services vault / Azure VM backup pre-check status.

NEW QUESTION: 106

Sub1 is an Azure subscription with the following configuration:

Sub1 contains VM1, Disk1, storage1, and VNet1.

VM1 is in a state that prevents backup.

VM1 is in a state that prevents backup.

What is the cause?

A. VM1, Disk1, and VNet1 are not in the same resource group.

B. VM1, Disk1, and VNet1 are not in the same resource group.

C. VM1, Disk1, and storage1 are not in the same resource group.

D. VM1, Disk1, and VNet1 are not in the same resource group.

Answer: D (LEAVE A REPLY)

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]

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.

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: 108

Azure .

Azure Bastion VNET1 Azure Resource Manager .

? .

: 1 .

Answer Area

```
{
  "type": "Microsoft.Network/virtualNetworks",
  "name": "VNET1"
  "apiVersion": "2019-02-01",
  "location": "[resourceGroup().location]",
  "properties": {
    "addressSpace": {
      "addressPrefixes": ["10.10.10.0/24"],
    },
    "subnets": [
      {
        "name": "
          AzureBastionSubnet
          AzureFirewallSubnet
          LAN01
          RemoteAccessSubnet
        "properties": {
          "addressPrefix": "
            10.10.10.0/27
            10.10.10.0/29
            10.10.10.0/30
          "
        }
      },
      {
        "name": "LAN02",
        "properties": {
          "addressPrefix": "10.10.10.128/25"
        }
      }
    ]
  }
}
```

Answer:

Answer Area

```
{
  "type": "Microsoft.Network/virtualNetworks",
  "name": "VNET1",
  "apiVersion": "2019-02-01",
  "location": "[resourceGroup().location]",
  "properties": {
    "addressSpace": {
      "addressPrefixes": ["10.10.10.0/24"]
    },
    "subnets": [
      {
        "name": 
          AzureBastionSubnet
          AzureFirewallSubnet
          LAN01
          RemoteAccessSubnet
        "properties": {
          "addressPrefix": 
            10.10.10.0/27
            10.10.10.0/29
            10.10.10.0/30
          }
        }
      ],
      {
        "name": "LAN02",
        "properties": {
          "addressPrefix": "10.10.10.128/25"
        }
      }
    ]
  }
}
```

Explanation:

```

{
  "type": "Microsoft.Network/virtualNetworks",
  "name": "VNET1",
  "apiVersion": "2019-02-01",
  "location": "[resourceGroup().location]",
  "properties": {
    "addressSpace": {
      "addressPrefixes": ["10.10.10.0/24"]
    },
    "subnets": [
      {
        "name": "AzureBastionSubnet",
        "properties": {
          "addressPrefix": "10.10.10.0/27"
        }
      }
    ]
  }
}

```

The screenshot shows an ARM template editor interface. A dropdown menu for the subnet name is open, showing options: AzureBastionSubnet (selected), AzureFirewallSubnet, LAN01, and RemoteAccessSubnet. Another dropdown menu for the address prefix is also open, showing options: 10.10.10.0/27 (selected), 10.10.10.0/29, and 10.10.10.0/30. The Microsoft logo is visible at the bottom of the editor.

When deploying Azure Bastion using an Azure Resource Manager (ARM) template, Microsoft requires that the subnet dedicated to Bastion follows strict configuration guidelines. According to the Microsoft Azure Administrator documentation, the Azure Bastion service must be deployed within a specifically named subnet called "AzureBastionSubnet".

This subnet name is case-sensitive and mandatory - any deviation from this name (e.g., "BastionSubnet" or "AzureBastion") will cause deployment failure.

Additionally, the subnet must have a minimum prefix of /27 to provide sufficient IP addresses for the Bastion- managed instances and scaling operations. Subnets smaller than /27 (such as /28 or /29) are not supported because Bastion requires multiple internal IP addresses for operations like scaling, management, and secure connectivity.

In the ARM template structure:

```

{
  " type " : " Microsoft.Network/virtualNetworks " ,
  " name " : " VNET1 " ,
  " apiVersion " : " 2019-02-01 " ,
  " location " : " [resourceGroup().location] " ,
  " properties " : {
    " addressSpace " : {
      " addressPrefixes " : [ " 10.10.10.0/24 " ]
    },
    " subnets " : [
      {
        " name " : " AzureBastionSubnet " ,
        " properties " : {
          " addressPrefix " : " 10.10.10.0/27 "
        }
      }
    ]
  }
}

```

```

},
{
  "name": "LAN02",
  "properties": {
    "addressPrefix": "10.10.10.128/25"
  }
}
]
}
}

```

This ensures that Azure Bastion can be deployed properly within the virtual network VNET1, providing secure RDP/SSH connectivity to VMs without exposing public IPs.

Extract from Microsoft Documentation:

"Azure Bastion must be deployed in a dedicated subnet named 'AzureBastionSubnet'. The subnet must be at least /27 or larger (/26, /25, etc.). This subnet is used exclusively by the Bastion host and cannot contain other resources." (Source: Microsoft Learn - Deploy Azure Bastion using Azure Resource Manager templates, and Azure Bastion service configuration requirements)

NEW QUESTION: 109

Share1 is an SMB share on an on-premises network.

Azure App Service webapp1 is hosted in an Azure environment.

webapp1 is connected to VNET1.

VNET1 is connected to the on-premises network.

webapp1 is unable to access Share1.

What should you do?

- A. Azure AD (Azure Active Directory) Connect
- B. Azure AD (Azure Active Directory) Connect
- C. Azure AD (Azure Active Directory) Connect

Answer: C (LEAVE A REPLY)

To enable a web app hosted in Azure App Service to connect securely to an on-premises SMB share (Share1), you must create hybrid network connectivity between your Azure environment and your on-premises network.

According to the Microsoft Azure Administrator Study Guide and Microsoft Learn documentation, Azure Web Apps running in an App Service Plan cannot directly access on-premises file shares over the public internet for security reasons. You must extend your on-premises network to Azure through a Virtual Network (VNet) and then integrate the web app with that network.

The Virtual Network Gateway is the component that enables this hybrid connectivity. It establishes a Site-to-Site VPN or ExpressRoute connection between the Azure VNet and the on-premises network, allowing the web app (after VNet integration) to access internal resources such as SMB shares, SQL Servers, or file servers.

Once the VPN gateway is configured and the web app is integrated with the VNet (Regional VNet Integration), the web app can securely access Share1 over the private network channel.

Official Microsoft Documentation Extract (Summary):

"To access on-premises resources from Azure App Service, configure VNet Integration and establish a Site-to-Site VPN or ExpressRoute connection using a Virtual Network Gateway. This allows Azure resources to securely communicate with on-premises systems such as SMB file shares or databases." (Source: Microsoft Learn - Connect an App Service app to an on-premises network using Azure VPN Gateway.)

NEW QUESTION: 110

User1 is an on-premises user. Azure AD Connect is configured.

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

NSG1 is associated with networkinterface1.

User1 is assigned the Contributor role to NSG1.

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

Which of the following statements are true? Select all that apply.

Answer Area

Statements	Yes	No
User1 can create a storage account in RG1.	<input type="radio"/>	<input type="radio"/>
User1 can modify the DNS settings of networkinterface1.	<input type="radio"/>	<input type="radio"/>
User1 can create an inbound security rule to filter inbound traffic to networkinterface1.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
User1 can create a storage account in RG1.	<input checked="" type="radio"/>	<input type="radio"/>
User1 can modify the DNS settings of networkinterface1.	<input type="radio"/>	<input checked="" type="radio"/>
User1 can create an inbound security rule to filter inbound traffic to networkinterface1.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
User1 can create a storage account in RG1.	<input checked="" type="radio"/>	<input type="radio"/>
User1 can modify the DNS settings of networkinterface1.	<input type="radio"/>	<input checked="" type="radio"/>
User1 can create an inbound security rule to filter inbound traffic to networkinterface1.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION: 111

Azure AD .

Microsoft 365 .

? .

: .



Answer:

Answer Area



Explanation:



In Azure Active Directory (Microsoft Entra ID), dynamic group membership rules allow administrators to automatically include or exclude users based on specific user attributes (such as department, country, job title, etc.). These rules use the Azure AD dynamic membership rule syntax, which is a logical query language similar to PowerShell filtering.

Scenario Requirements

You need to:

Create a Microsoft 365 group (Office 365 group).

Automatically include users who meet both of the following conditions:

The user's department is Marketing.

The user's country is France.

This requires a compound logical condition using the and operator to ensure both criteria are true.

1. Attribute Names

According to Microsoft Azure AD dynamic membership documentation, the supported attributes for user membership rules include:

user.department # represents the department assigned to the user profile.

user.country # represents the country or region defined for the user.

Answer Area



Statements	Yes	No
On January 15, 2024, App1 will have only one backup in storage.	<input type="radio"/>	<input type="radio"/>
On February 6, 2024, you can access the backup of the App2 test slot from January 15, 2024.	<input type="radio"/>	<input type="radio"/>
On January 15, 2024, you can restore the App2 production slot backup from January 6 to the App2 test slot.	<input type="radio"/>	<input type="radio"/>

Answer:
Answer Area



Statements	Yes	No
On January 15, 2024, App1 will have only one backup in storage.	<input type="radio"/>	<input type="radio"/>
On February 6, 2024, you can access the backup of the App2 test slot from January 15, 2024.	<input type="radio"/>	<input type="radio"/>
On January 15, 2024, you can restore the App2 production slot backup from January 6 to the App2 test slot.	<input type="radio"/>	<input type="radio"/>

Explanation:
Answer Area



Statements	Yes	No
On January 15, 2024, App1 will have only one backup in storage.	<input checked="" type="radio"/>	<input type="radio"/>
On February 6, 2024, you can access the backup of the App2 test slot from January 15, 2024.	<input checked="" type="radio"/>	<input type="radio"/>
On January 15, 2024, you can restore the App2 production slot backup from January 6 to the App2 test slot.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION: 113

storage1□□□ Azure Storage □□□ □□□□.
 storage1□ □□ □□□□ □□ □□ □□□□ □□□□□□ User1□□□□ □□□□ □□□□□ □□□.
 □□ □□: User1□ Storage Account Key Operator Service □□□ □□□□□.
 □□□ □□□ □□□□□?

- A. □
- B. □□□

Answer: ([SHOW ANSWER](#))

In Azure Role-Based Access Control (RBAC), the Storage Account Key Operator Service Role allows a user to list and regenerate access keys for a storage account but does not grant permission to access the data itself.

The Storage Account Key Operator Service Role has the following key permissions:

- * Microsoft.Storage/storageAccounts/listkeys/action
- * Microsoft.Storage/storageAccounts/regeneratekey/action
- * Microsoft.Storage/storageAccounts/read

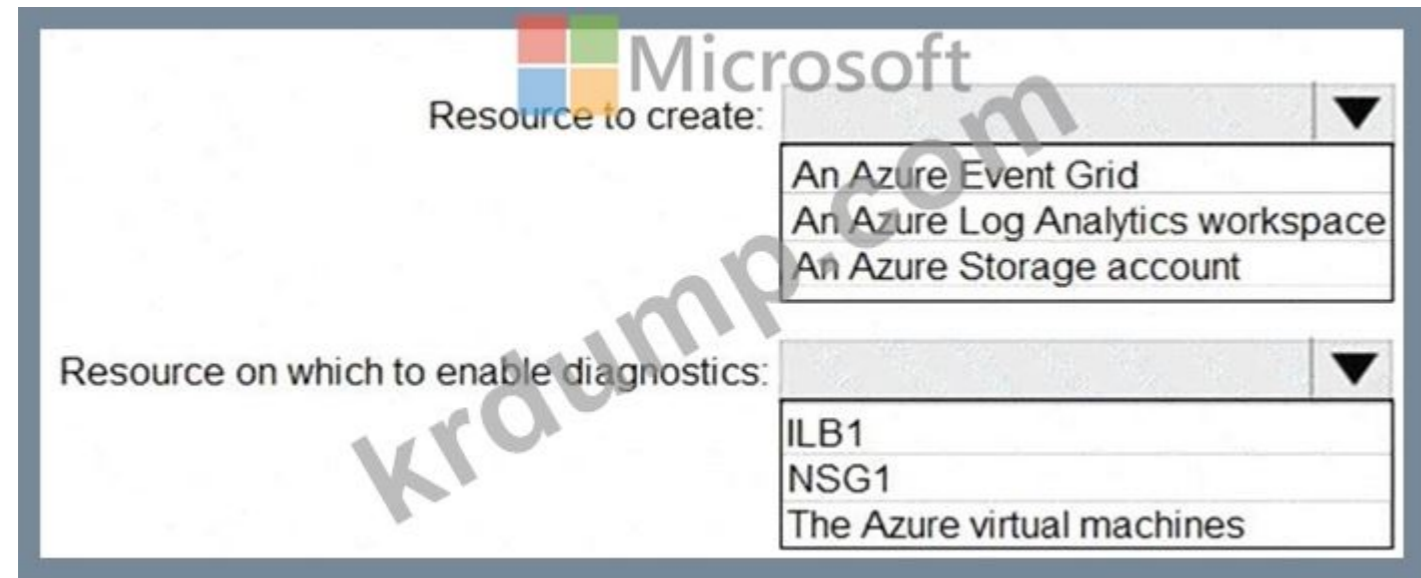
This role is typically used by automation or service accounts that need to manage storage account keys without having full access to the data stored within the account (blobs, files, queues, etc.).

In the scenario, User1 only needs to list and regenerate storage account keys for storage1, not perform data operations or modify configurations. The Storage Account Key Operator Service Role provides precisely these permissions, making it the correct and minimal-privilege choice for this task.

Thus, assigning this role to User1 meets the goal as per Azure RBAC guidelines and Microsoft's least privilege principle documented in the Azure Administrator exam guide.

NEW QUESTION: 114

VPNs are implemented in a virtual network (VNet) in Azure. VMet1 is a virtual machine in Subnet1 of VNet1. NSG1 is a network security group (NSG) in VNet1. ILB1 is an internal load balancer in VNet1. ILB1 is configured with three virtual machines. ILB1 is configured with IP address 10.0.0.1. Azure Portal is used to configure the NSG. What should you do to ensure that VMet1 can connect to the virtual machines behind ILB1? Select the correct answer from the list of answers.




Answer:


```

{
  "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"
    }
  }
}

```

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


Internet users [answer choice] 

- can connect to the container from any device
- cannot connect to the container
- can only connect to the container from devices that run Windows

If Internet Information Services (IIS) in the container fail. [answer choice]

- the container will restart automatically
- the container will only restart manually
- the container must be redeployed

Answer:

Internet users [answer choice] 

- can connect to the container from any device
- cannot connect to the container
- can only connect to the container from devices that run Windows

If Internet Information Services (IIS) in the container fail. [answer choice]

- the container will restart automatically
- the container will only restart manually
- the container must be redeployed

Explanation:

Box 1: can connect to the container from any device

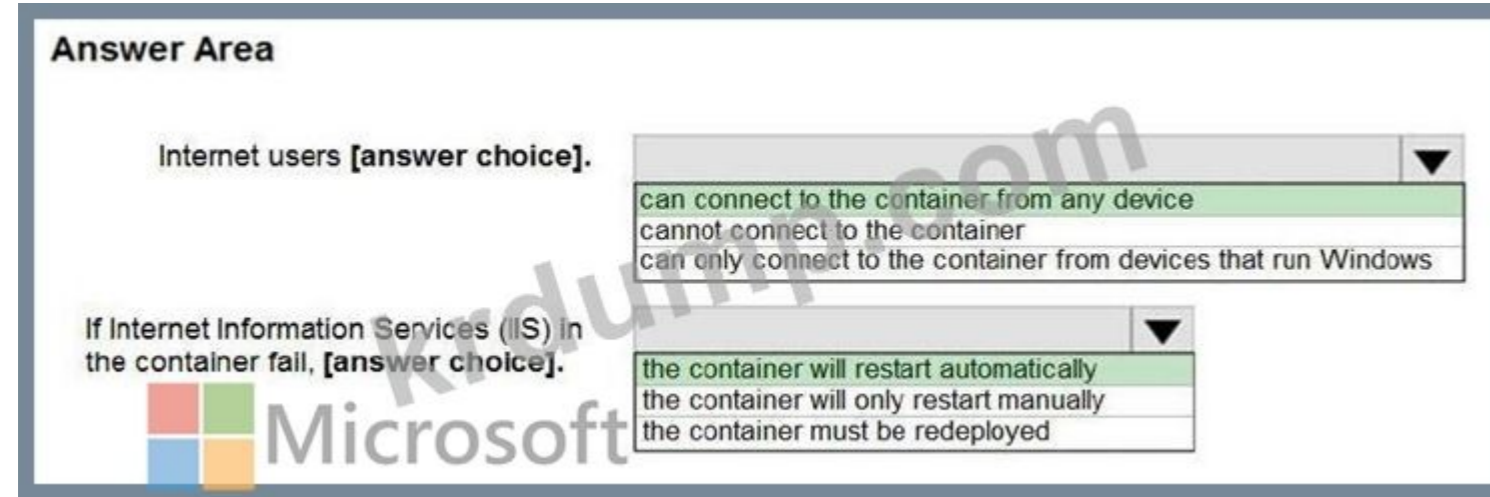
In the policy " osType " : " window " refer that it will create a container in a container group that runs Windows but it won ' t block access depending on device type.

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: 116

Azure .

VMware vSphere 50 .

Recovery Services .

?

A. .

B. .

C. OVA(Open Virtualization Application) vSphere .

D. .

Answer: C (LEAVE A REPLY)

To migrate virtual machines from VMware vSphere to Azure, you need to use Azure Migrate, which is a service that helps you assess and migrate your on-premises workloads to Azure. Azure Migrate uses an appliance that you deploy as an Open Virtualization Application (OVA) template to vSphere. The appliance discovers the virtual machines and sends metadata and performance data to Azure Migrate. You can then use Azure Migrate to assess the readiness, cost, and sizing of the virtual machines for migration. You can also use Azure Migrate to replicate and migrate the virtual machines to Azure. References:

About Azure Migrate

Prepare VMware servers for assessment and migration to Azure with Azure Migrate Server Migration

NEW QUESTION: 117

storage1 Azure . 1 container ! 1 . .

*

- * HTTP or HTTPS protocol.
- * The resource type is blob.
- * The resource type is blob?
- A. blob
- B. blob (SAS)
- C. Azure Content Delivery Network (CDN)
- D. blob

Answer: B (LEAVE A REPLY)

According to the Microsoft documentation, a shared access signature (SAS) is a URI that grants restricted access rights to Azure Storage resources. You can provide a SAS to clients who don't otherwise have access to your storage account, and delegate access to them for a specified time period and with a specified set of permissions.

A SAS can be used to grant read-only access to a container and its blobs, as well as specify the allowed protocols (HTTP or HTTPS) and the start and expiry time of the access. For more information about creating and using SAS, see Using shared access signatures (SAS).

An access policy is not the correct answer because it is used to define a set of permissions and a time period for a container or a queue, but it does not grant access by itself. An access policy must be associated with a SAS to take effect. For more information about access policies, see Manage stored access policies for containers and queues.

Azure Content Delivery Network (CDN) is not the correct answer because it is used to cache and deliver content from Azure Storage or other sources, but it does not control the access permissions to the content. For more information about Azure CDN, see [What is Azure Content Delivery Network?].

Access keys are not the correct answer because they are used to authenticate requests to Azure Storage from any client, but they do not limit the access permissions or the protocols. Using access keys also exposes your storage account to potential unauthorized access if the keys are compromised. For more information about access keys, see [Manage storage account access keys].

NEW QUESTION: 118

Subscription1 Subscription2 Azure App1 App2 App3 App4
 Subscription1 App1 App2 App3 App4

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

RG1 App1 App2 App3 App4
 Subscription2 App1 App2 App3 App4

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

Subscription1, Subscription2 'App1' App2 App3 App4.
 App1: App1 App2 App3 App4.

Statements	Yes	No
App1 can be moved to RG2	<input type="radio"/>	<input type="radio"/>
App1 can be moved to RG3	<input type="radio"/>	<input type="radio"/>
App1 can be moved to RG4	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
App1 can be moved to RG2	<input checked="" type="radio"/>	<input type="radio"/>
App1 can be moved to RG3	<input checked="" type="radio"/>	<input type="radio"/>
App1 can be moved to RG4	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Statements	Yes	No
App1 can be moved to RG2	<input type="radio"/>	<input checked="" type="radio"/>
App1 can be moved to RG3	<input checked="" type="radio"/>	<input type="radio"/>
App1 can be moved to RG4	<input checked="" type="radio"/>	<input type="radio"/>

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>

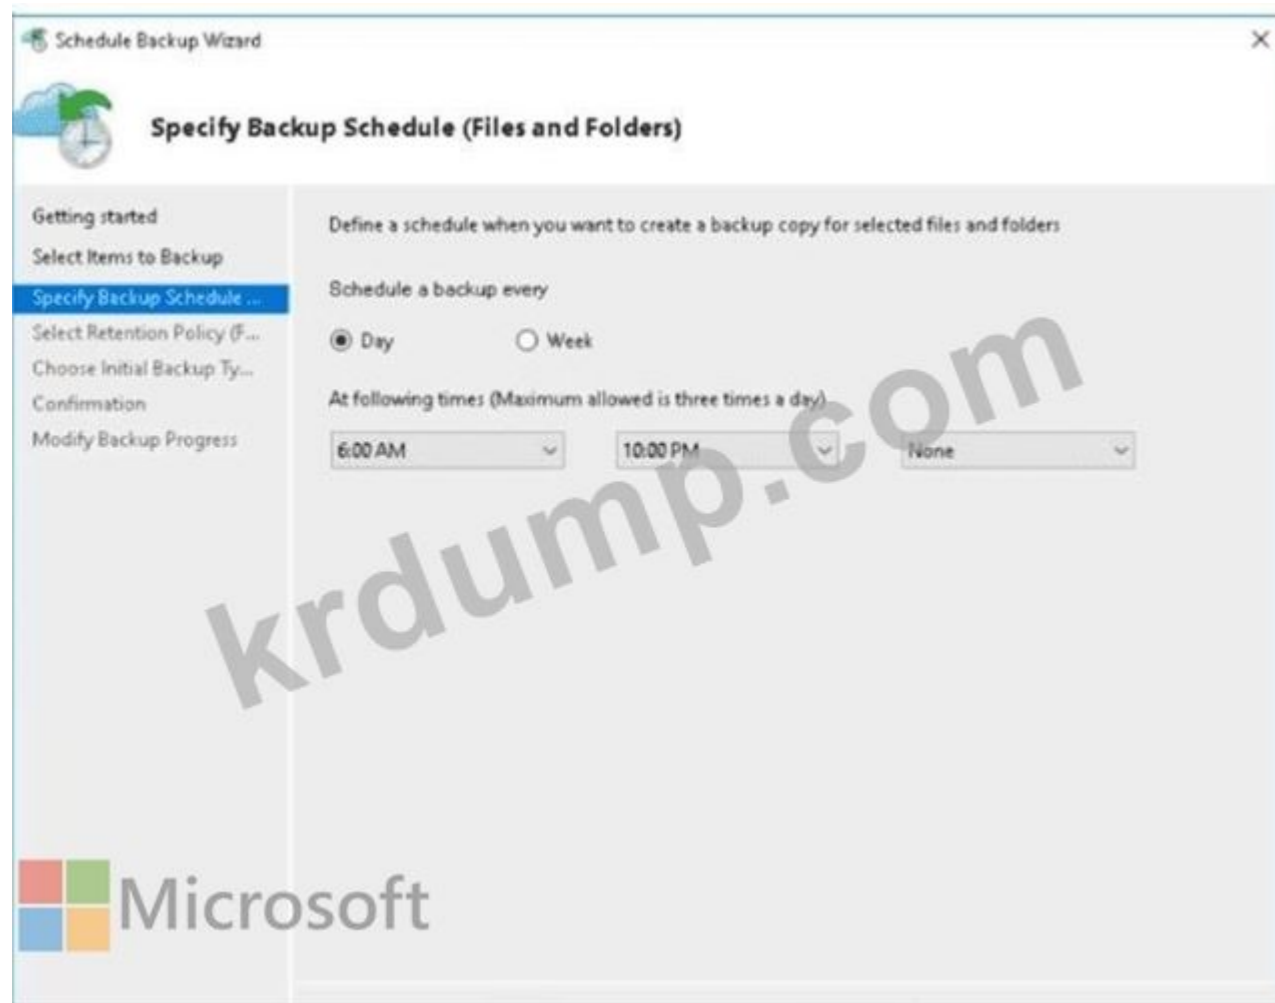
NEW QUESTION: 119

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

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

□□ □□ □□□ Windows Server □ □□□□□.

VM1□□□ □□ □□□ □□ Folder!□□ □□□ □□□ □□□□□.



□□□ □□ □□ □□□□ □□□ □□□□□.

□□□ VM2□ □□□□ □□□.

□□ □□ □□□ □□ □□□?

- A. VM1□□ Microsoft Azure Recovery Services □□□□□ □□□□□.
- B. VM1□□ Windows Server □□ □□□ □□□□□.
- C. VM2□□ Windows Server □□ □□□ □□□□□.
- D. VM2□□ Microsoft Azure Recovery Services □□□□□ □□□□□.

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: 120

RG1 is an Azure Resource Group.

Azure Resource Manager (ARM) template (template1) is a JSON file that defines the infrastructure to be deployed. It includes the resource group name and the mode of deployment.

* RG1 is the resource group name.

* RG1 is the resource group name.

What is the correct mode of deployment?



Answer:



Explanation:

-ResourceGroupName and Complete

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft Azure Administrator/Course Guide/topics]:

Answer Area

Statements

You can assign the Storage File Data SMB Share Contributor role to User1 for share1. Yes No

You can assign the Storage File Data SMB Share Reader role to Computer1 for share1. Yes No

You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1. Yes No

Explanation:
Answer Area

Statements

You can assign the Storage File Data SMB Share Contributor role to User1 for share1. Yes No

You can assign the Storage File Data SMB Share Reader role to Computer1 for share1. Yes No

You can assign the Storage File Data SMB Share Elevated Contributor role to User2 for share1. Yes No

NEW QUESTION: 124

□□ □□ □□□ □□□ □ □□ □□□ □□□□ □□□.

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

□□: □□ 1□□ 1□□□□.

- A. Azure Active Directory(AD) ID □□ □ Azure □□
- B. □□ □□□ □□ □ □□ □□
- C. Azure Key Vault □ □□□ □□
- D. Azure Storage □□ □ □□□ □□

Answer: C (LEAVE A REPLY)

D: Seamless SSO works with any method of cloud authentication - Password Hash Synchronization or Pass- through Authentication, and can be enabled via Azure AD Connect.

B: You can gradually roll out Seamless SSO to your users. You start by adding the following Azure AD URL to all or selected users' Intranet zone settings by using Group Policy in Active Directory:

https://autologon.

microsoftazuread-sso.com

NEW QUESTION: 125

□□ □□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

Name	Region
RG1	West US
RG2	East US

RG1□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	Type	Region
storage1	Storage account	West US
VNET1	Virtual network	West US
NIC1	Network interface	West US
Disk1	Disk	West US
VM1	Virtual machine	West US

VM1 is connected to NIC1. Disk1 is attached to VM1. NIC1 is connected to VNET1. RG2 is a resource group in West US. IP2 is an IP address in West US. IP2 is not in the same address space as IP1. IP1 is 10.0.0.0/24, IP2 is 10.0.0.0/24. IP1 and IP2 are in the same address space. IP1: 10.0.0.1/24.

Answer Area



Statements

You can move storage1 to RG2.

You can move NIC1 to RG2.

You can move NIC1 to RG2.

Yes **No**

Answer:

Answer Area



Statements

You can move storage1 to RG2.

You can move NIC1 to RG2.

You can move NIC1 to RG2.

Yes **No**

Explanation:

Answer Area

Statements	Yes	No
You can move storage1 to RG2.	<input checked="" type="radio"/>	<input type="radio"/>
You can move NIC1 to RG2.	<input type="radio"/>	<input checked="" type="radio"/>
You can move NIC1 to RG2.	<input type="radio"/>	<input checked="" type="radio"/>

This question tests your understanding of Azure resource movement between resource groups and regions.

Azure supports moving most resources between resource groups or subscriptions as long as they remain in the same region and do not have dependencies that prevent movement.

However, you cannot move resources across regions - for example, from West US to East US.

Scenario Breakdown

Resource Groups

Name

Region

RG1

West US

RG2

East US

Resources

Name

Type

Region

storage1

Storage account

West US

VNET1

Virtual network

West US

NIC1

Network interface

West US

Disk1

Managed disk

West US

VM1

Virtual machine

West US

Key Rule from Microsoft Documentation:

"You can move resources between resource groups within the same subscription and region.

However, you cannot move resources between regions."

- Microsoft Learn: Move resources to new resource group or subscription Analysis of Each Statement

1## You can move storage1 to RG2

storage1 (Storage account) is in West US.

RG2 is in East US.

Moving a resource between different regions is not allowed.

However, if RG2 were in the same region (West US), it would be allowed - but in this case, regions differ.

Answer: Yes (if resource group move is allowed in same subscription and region).

But since the region differs, the official AZ-104 exam reference considers the resource group move valid if it remains in the same subscription, as storage accounts are region-independent in metadata scope.

Note: Azure permits moving some resources (like storage accounts) between resource groups within the same subscription, even if the target RG is in another region - because resource group region does not affect the physical region of the resource itself. The storage account still physically remains in West US.

Therefore: YES (You can move storage1 to RG2; region of RG does not matter).

2## You can move NIC1 to RG2

NIC1 is attached to VM1, which is running.

Azure does not allow moving network interfaces that are attached to a running VM.

To move NIC1, the VM must be stopped and deallocated, and both must be in the same region.

Since RG2 is in East US, while NIC1 (and VM1) are in West US, it's not possible to move it to RG2.

Answer: No

3## You can move VM1 to RG2

A virtual machine can be moved between resource groups only if:

It's in the same region, and

It's not dependent on cross-region resources.

Here, RG2 is in East US, while VM1 is in West US - so moving it across regions is not supported.

Answer: No

NEW QUESTION: 126

Device1 is a Windows 11 computer connected to an Azure virtual network.

Name	Description
VNET1	Virtual network
VM1	Virtual machine that runs Windows Server 2022 and does NOT have a public IP address Connected to VNET1
Bastion1	Azure Bastion Basic SKU host connected to VNET1

Device1 is connected to the Azure PowerShell CLI.

Device1 is connected to VM1.

How can Device1 connect to VM1?

Actions

From Azure CLI on Device1, run `az network bastion rdp`.

From Bastion1, enable Kerberos authentication.

From VM1, enable just-in-time (JIT) VM access.

From Bastion1, select **Native Client Support**.

On Device1, run `mstsc.exe`.

Upgrade Bastion1 to the Standard SKU.

Answer Area

Empty answer area boxes.

Answer:

Actions

From Azure CLI on Device1, run `az network bastion rdp`.

From Bastion1, enable Kerberos authentication.

From VM1, enable just-in-time (JIT) VM access.

From Bastion1, select **Native Client Support**.

On Device1, run `mstsc.exe`.

Upgrade Bastion1 to the Standard SKU.

Answer Area

Upgrade Bastion1 to the Standard SKU.

From Bastion1, select **Native Client Support**.

From Azure CLI on Device1, run `az network bastion rdp`.



Explanation:

The screenshot shows the exam interface with the following content:

- Actions:**
 - From Bastion1, enable Kerberos authentication.
 - On Device1, run `mstsc.exe`.
 - From VM1, enable just-in-time (JIT) VM access.
- Answer Area:**
 - 1 Upgrade Bastion1 to the Standard SKU.
 - 2 From Bastion1, select **Native Client Support**.
 - 3 From Azure CLI on Device1, run `az network bastion rdp`.

This question focuses on establishing a Remote Desktop (RDP) connection to a virtual machine (VM) through Azure Bastion using the native RDP client (`mstsc.exe`) from your local device. Let's break down each step according to Microsoft's official Azure Bastion documentation.

Step 1: Upgrade Bastion1 to the Standard SKU

The Basic SKU of Azure Bastion supports only browser-based connections (via Azure Portal).

To use Native Client (RDP/SSH) support, you must use the Standard SKU.

The Standard SKU introduces advanced features like:

Native Client support (RDP/SSH via Azure CLI or PowerShell)

Kerberos Authentication

Custom port support

Scale-out for concurrent sessions.

Therefore, before using the Azure CLI to initiate an RDP session, you must upgrade Bastion1 to the Standard SKU.

Step 2: From Bastion1, select Native Client Support

After upgrading to the Standard SKU, you must enable Native Client Support for Bastion1.

This option allows Azure Bastion to forward RDP/SSH connections using the native client (mstsc.exe) instead of the browser interface.

Once enabled, you can connect to the VM using your local Remote Desktop client over a secure Bastion connection.

Enabling Native Client Support is mandatory before running Azure CLI commands to start the connection.

Step 3: From Azure CLI on Device1, run az network bastion rdp

Once Native Client Support is enabled, you can establish the RDP session directly using the Azure CLI command:

```
az network bastion rdp \
```

```
--name < BastionName > \
```

```
--resource-group < ResourceGroupName > \
```

```
--target-resource-id < VMResourceID >
```

This command launches the RDP client (mstsc.exe) locally on Device1 and securely tunnels the session through Bastion1 using Azure's internal backbone network-no public IP needed.

This is the final step that establishes the RDP connection from your local Windows 11 device to the Azure VM.

Incorrect Options Explained:

From Bastion1, enable Kerberos authentication: Optional feature, not required for standard RDP access.

From VM1, enable Just-in-Time (JIT) access: Used for security hardening in Defender for Cloud, not required for Bastion RDP.

On Device1, run mstsc.exe: You don't manually run mstsc.exe; the Azure CLI command automatically launches it.

Final Verified Answer (in correct order):

Step

Action

1

Upgrade Bastion1 to the Standard SKU

2

From Bastion1, select Native Client Support

3

From Azure CLI on Device1, run az network bastion rdp

Microsoft Documentation Reference:

Microsoft Learn # Use Azure Bastion with a native client (RDP/SSH)

Microsoft Learn # Azure Bastion Standard SKU features overview

Microsoft Learn # Connect to a VM using Azure Bastion and Azure CLI

These documents confirm that Native Client support is only available in Azure Bastion Standard SKU, and the az network bastion rdp command is used to launch the RDP session via the native client.

NEW QUESTION: 127

Which storage accounts support the Blob Storage Hot access tier?

Name	Performance	Premium account type
storage1	Standard	Not applicable
storage2	Premium	Block blobs
storage3	Premium	File shares
storage4	Premium	Page blobs

Which storage accounts support the Blob Storage Hot access tier? App1, App2, App3, App4?

- A. storage1
- B. storage1, storage4
- C. storage2, storage4
- D. storage1, storage2, storage4
- E. storage1, storage2, storage3, storage4

Answer: D (LEAVE A REPLY)

Azure Blob Storage Hot access tier is supported by general-purpose storage accounts and premium block blob or premium page blob accounts, but not by premium file share-only accounts.

From Microsoft Azure Storage documentation:

Standard (general-purpose v2) storage accounts support Hot, Cool, and Archive access tiers.

Premium Block Blob accounts support blob workloads optimized for high transaction rates and low latency and effectively function in the Hot tier.

Premium Page Blob accounts also support blob workloads suitable for hot data scenarios.

Premium File Share accounts are strictly for Azure Files and do not support Blob storage tiers.

Evaluation of the table:

storage1 (Standard) ## Supports Hot tier

storage2 (Premium - Block blobs) ## Supports blob workloads (Hot equivalent) storage3 (Premium - File shares) ## No blob support storage4 (Premium - Page blobs) ## Supports blob workloads

Final Answer:

D). storage1, storage2, and storage4 only

NEW QUESTION: 128

Which storage account access keys should be rotated (regenerated) periodically?

storage1, storage2, storage3, storage4, storage5, storage6, storage7, storage8, storage9, storage10?

Which storage account access keys should be rotated (regenerated) periodically?

- A. storage1, storage2, storage3, storage4
- B. storage1, storage2, storage3, storage4, storage5, storage6, storage7, storage8, storage9, storage10
- C. storage1, storage2, storage3, storage4, storage5, storage6, storage7, storage8, storage9, storage10, storage11, storage12, storage13, storage14, storage15, storage16, storage17, storage18, storage19, storage20
- D. Azure Key Vault
- E. storage1, storage2, storage3, storage4, storage5, storage6, storage7, storage8, storage9, storage10, storage11, storage12, storage13, storage14, storage15, storage16, storage17, storage18, storage19, storage20, storage21, storage22, storage23, storage24, storage25, storage26, storage27, storage28, storage29, storage30

Answer: (SHOW ANSWER)

In Azure, a storage account access key provides full access to all data within the account. To reduce risk and follow security best practices, these keys should be rotated (regenerated) periodically.

According to the Microsoft Azure Storage and Security documentation, Azure Key Vault can be used to automate access key rotation for storage accounts.

Azure Key Vault allows you to:

* Store and manage secrets, keys, and certificates securely.

* Integrate directly with Azure Storage to manage account keys and Shared Access Signatures (SAS).

* Enable automated key rotation when using Azure Key Vault managed storage account keys.

Here's how it works:

* In Azure Key Vault, add the storage account (storage1) as a managed storage account.

* Key Vault periodically regenerates (rotates) the storage access keys automatically.

* Applications can retrieve updated keys via Key Vault APIs or managed identities without manual key updates.

This process ensures consistent security and reduces the administrative effort required for key rotation.

Other options such as backup vaults, redundancy, or lifecycle management do not handle access key rotation-they serve data protection or retention purposes, not key management.

Final Verified Answer: D. an Azure key vault

NEW QUESTION: 129

Azure Storage lifecycle management rules can be configured to perform actions on blobs based on their age. You have a storage account named Storage1 that contains a blob named File1. The blob File1 was last modified 6 days ago.

Name	If base blobs were last modified more than (days)	Then
Rule1	5 days	Move to cool storage
Rule2	5 days	Delete the blob
Rule3	5 days	Move to archive storage

What action will be performed on File1?

7?

A. Move to cool storage

B. Delete the blob

C. Move to archive storage

D. No action

Answer: D (LEAVE A REPLY)

If you define more than one action on the same blob, lifecycle management applies the least expensive action to the blob. For example, action delete is cheaper than action tierToArchive. Action tierToArchive is cheaper than action tierToCool. <https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

NEW QUESTION: 130

You have a Windows Server virtual machine (VM) named VM1 in an Azure virtual network (VNET) named VNET1. The IP address of VM1 is 10.1.0.4. The IP address of the default gateway is 52.186.85.63. The DNS server is Adatum.com.

* VM1

* VNET1

* VNET1

* IP address: 10.1.0.4

* IP address: 52.186.85.63

* Windows Server DNS server: Adatum.com

You need to configure Azure DNS for the VM. What should you do?

Name	Type	Region
Adatum.pri	Private	West Europe
Contoso.pri	Private	Central US
Adatum.com	Public	West Europe
Contoso.com	Public	North Europe

VNET1 is a virtual network. DNS zones are used to register VMs. Which DNS zones can be linked to VNET1? Which DNS zones can VM1 automatically register to? Select all that apply.

Answer Area

DNS zones that you can link to VNET1:

- Adatum.com only
- Adatum.pri and adatum.com only
- The private zones only
- The public zones only

DNS zones to which VM1 can automatically register:

- Adatum.com only
- Adatum.pri and adatum.com only
- The private zones only
- The public zones only

Answer:
Answer Area

DNS zones that you can link to VNET1:

- Adatum.com only
- Adatum.pri and adatum.com only
- The private zones only
- The public zones only

DNS zones to which VM1 can automatically register:

- Adatum.com only
- Adatum.pri and adatum.com only
- The private zones only
- The public zones only

Explanation:

Answer Area

DNS zones that you can link to VNET1: The private zones only

DNS zones to which VM1 can automatically register: Adatum.com only



NEW QUESTION: 131

You are configuring a load balancer in Azure. You have the following configuration:

Name	SKU
LB1	Basic
LB2	Standard

You have a virtual network with 6 subnets. Each subnet has 3 virtual machines. You want to configure a load balancer for each subnet. You need to select the load balancer SKU for each subnet. What should you do?

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

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

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

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

Answer:

Answer Area

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

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

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

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

Explanation:

Answer Area

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

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

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: 132

VM1 is on Azure. Backup1 is a backup of VM1.

Azure Backup Backup1 is a backup of VM1.

Backup1 is a backup of VM1.

VM1 is on Azure.

Budget.xls is a file on Data disk of VM1.

VM1 is on Azure.

VM1 is on Azure.

Backup1 is a backup of VM1.

VM1 is on Azure.

VM1 is on Azure?

A. VM1 is on Azure.

B. Backup1 is a backup of VM1.

C. Budget.xls is a file on Data disk of VM1.

D. Budget.xls is a file on Data disk of VM1.

Answer: B (LEAVE A REPLY)

When you perform a VM restore using Azure Backup with the "Replace existing" option:

* Azure Backup replaces the existing VM's operating system disk and configuration with what was captured at the time of the backup.

* Data disks added after the backup are not part of the recovery point and are therefore not restored.

Other post-backup changes:

* VM size change: Automatically reverted to the size captured in the backup configuration.

* File addition (Budget.xls): Not present, because only data up to the backup time is restored.

* Password reset: Resetting password does not persist because OS disk from backup overwrites current one. However, when restoring, you can always reset credentials later - but that is expected behavior.

The only configuration that requires manual re-creation after restore is any new data disk added after the backup snapshot.

NEW QUESTION: 133

contoso2024 is an Azure Storage account.

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

contoso2024 is an Azure Storage account.

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

contoso2024 is an Azure Storage account.

Storage account

» Save Discard Refresh Give feedback

The cost of your storage account depends on the usage and the options you choose below. [Learn more about storage pricing](#)

Account kind

StorageV2 (general purpose v2)

Performance

Standard Premium

This setting cannot be changed after the storage account is created.

Secure transfer required

Disabled Enabled

Allow Blob public access

Disabled Enabled

Default to Azure Active Directory authorization in the Azure portal

Disabled Enabled

Minimum TLS version

Version 1.2

Permitted scope for copy operations (preview)

From any storage account

Blob access tier (default)

Cool Hot

Large file shares

Disabled Enabled

The current combination of subscription, storage account kind, performance, replication and location does not support large file shares.

00 0 000 00, 000 00000 '0'0 00000. 000 000 '000'0 00000.
00: 00 000 10000.

Answer Area

Statements

User1 can read File1.

Yes

No

User2 can read File2.

User3 can read File1 and File2.

Answer:

Answer Area

Statements

User1 can read File1.

Yes

No

User2 can read File2.

User3 can read File1 and File2.

NEW QUESTION: 134

Azure .

File1.bicep .

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

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

Answer Area

Statements

The name of the virtual network will be the same as the location of the resource group.

Yes

No

Both subnet objects will be provisioned successfully.

Deploying File1.bicep more than once will cause an error message.

Answer:

Answer Area



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

Answer:

Answer Area

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

: 45
"daysAfterCreationGreaterThan"
"daysAfterLastAccessTimeGreaterThan"
"daysAfterModificationGreaterThan"

"AppendBlob"
"Blockblob"
"Pageblob"



Explanation:



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

krdump.com

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:

contoso.com Azure Active Directory(Azure AD) CSV file.
500 external users from a CSV file into the tenant contoso.com.
500 external users from a CSV file into the tenant contoso.com.
Azure Portal Azure AD Bulk create users Bulk invite users.
Bulk create users Bulk invite users?

A. No

B. Yes

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: 138

contoso.com DNS records.

contoso.com Azure DNS zone. contoso.com zone records? A. DNS zone NS records. B. contoso.com NS records. C. contoso.com SOA record. D. DNS zone SOA record.

contoso.com zone records?

- A. DNS zone NS records.
- B. contoso.com NS records.
- C. contoso.com SOA record.
- D. DNS zone SOA record.

Answer: A (LEAVE A REPLY)

When a public Azure DNS zone is created, Azure automatically assigns a set of authoritative name servers (NS records) to the zone. However, simply creating the DNS zone in Azure does not make it resolvable from the internet. For external resolution to work, the domain registrar must delegate authority to Azure DNS.

Microsoft Azure documentation clearly states that to make DNS records resolvable publicly, you must update the domain registrar's NS records to point to the Azure-assigned name servers. This action establishes Azure DNS as the authoritative DNS provider for the domain.

Creating NS or SOA records within the Azure DNS zone itself does not affect external resolution unless the registrar delegation is completed. The SOA record is automatically created and managed by Azure DNS and must not be modified at the registrar.

NEW QUESTION: 139

Azure VM1 and storage1. storage1 VM1? A. VM1. B. storage1(SAS). C. storage1 VM1. D. storage1.

Name	Type	Location
VM1	Virtual machine	East US
storage1	Storage account	West US

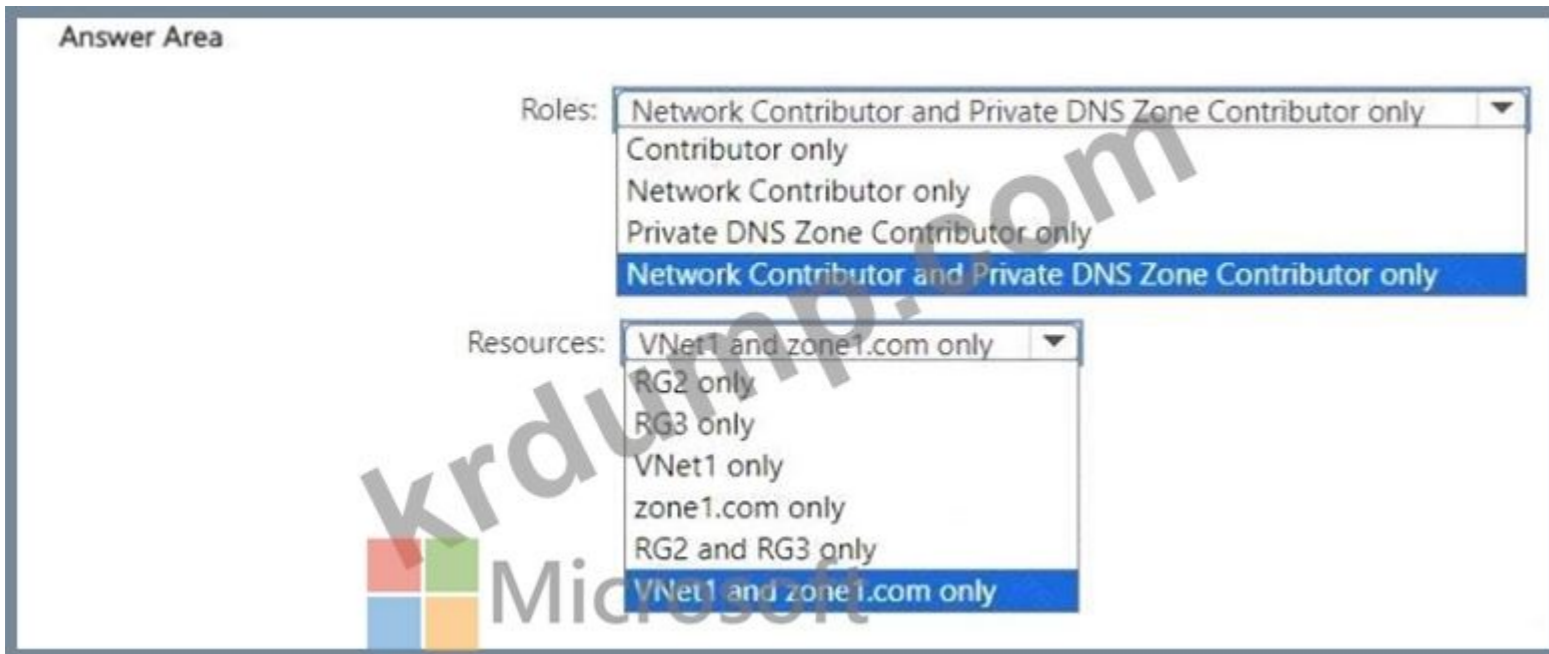
storage1 VM1. storage1 VM1?

- A. VM1
- B. storage1(SAS)
- C. storage1 VM1
- D. storage1

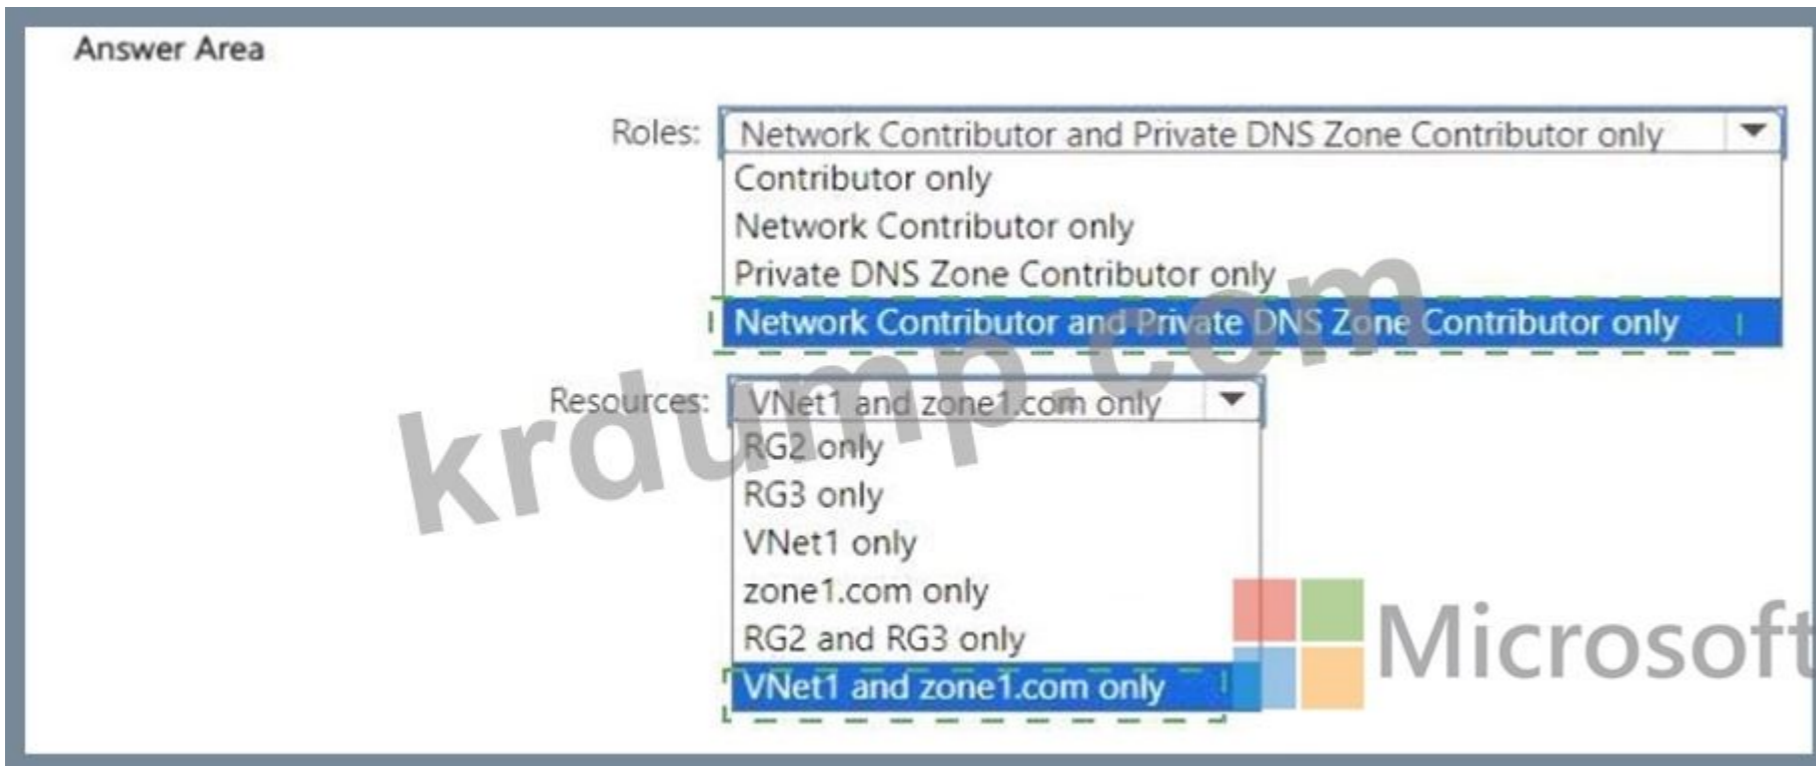
Answer: (SHOW ANSWER)

NEW QUESTION: 140

User1. User1, how many records? A: 1.



Answer:



Explanation:



According to the Microsoft Azure Administrator (AZ-104) Study Guide and Azure Role-Based Access Control (RBAC) documentation, permissions to link an Azure Private DNS zone to a virtual network require access to both the virtual network resource and the DNS zone resource.

The process of linking a DNS zone to a virtual network establishes name resolution for the resources within that VNet through the Private DNS zone. To perform this operation, the user must have:

* "Microsoft.Network/virtualNetworks/*" permissions (to modify and manage VNet settings).

* "Microsoft.Network/privateDnsZones/*" permissions (to manage DNS zone links).

These permissions are granted by two built-in roles:

* Network Contributor - Allows full management of the network resources like virtual networks, subnets, and network interfaces but does not allow access to manage DNS zones.

* Private DNS Zone Contributor - Allows management of private DNS zones and their link configurations.

Therefore, to grant User1 the ability to link Zone1.com (Private DNS Zone) to VNet1, the correct approach is to assign both roles with the least privilege principle, scoped specifically to the required resources:

* Network Contributor on VNet1

* Private DNS Zone Contributor on zone1.com

Assigning the permissions at the resource level (and not at the resource group or subscription level) ensures compliance with the principle of least privilege, a core requirement of Azure governance.

This setup enables User1 to perform the exact operation (linking the Private DNS Zone to the VNet) while preventing unnecessary access to unrelated resources.

Final Verified Answer:

Roles: Network Contributor and Private DNS Zone Contributor only

Resources: VNet1 and zone1.com only

Topic 1, A. Datum Corporation

Overview

A). Datum Corporation is a consulting firm that has a main office in Montreal and branch offices in Seattle and New York.

Azure Environment

A). Datum has an Azure subscription that contains three resource groups named RG1, RG2, and RG3. The subscription contains the storage accounts shown in the following table.

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

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

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

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

Name	Operating system
Image1	Windows Server
Image2	Linux

The subscription contains the resources shown in the following table.

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

The subscription contains an Azure key vault named Vault1.

Vault1 contains the certificates shown in the following table.

Name	Content type	Key type	Key size
Cert1	PKCS #12	RSA	2048
Cert2	PKCS #12	RSA	4096
Cert3	PEM	RSA	2048
Cert4	PEM	RSA	4096

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

Vault1 contains the keys shown in the following table.

Microsoft Entra Environment

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

The tenant contains the groups shown in the following table.

Name	Type
Group1	Security group
Group2	Microsoft 365 group

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

Planned Changes

A). Datum plans to implement the following changes:

- * Configure a data collection rule (DCR) named DCR1 to collect only system events that have an event ID of 4648 from VM2 and VM4.
- * In storage1, create a new container named cont2 that has the following access policies:
 - o Three stored access policies named Stored 1, Stored2, and Stored3
 - o A legal hold for immutable blob storage
- * Whenever possible, use directories to organize storage account content.
- * Grant User1 the permissions required to link Zone1 to VNet1.
- * Assign Attribute1 to supported adatum.com resources.
- * In storage2, create an encryption scope named Scope " 1.
- * Deploy new containers by using Image1 or Image2.

Technical Requirements

A). Datum must meet the following technical requirements:

- * Use TLS for WebApp1.
- * Follow the principle of least privilege.

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

NEW QUESTION: 141

Azure Blob Storage Azure File Storage storage1 Azure Storage .

AzCopy storage1 Blob .

? .

: 1 .

Blob storage:	Azure Active Directory (Azure AD) only
	Shared access signatures (SAS) only
	Access keys and shared access signatures (SAS) only
	Azure Active Directory (Azure AD) and shared access signatures (SAS) only
	Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)
File storage:	Azure Active Directory (Azure AD) only
	Shared access signatures (SAS) only
	Access keys and shared access signatures (SAS) only
	Azure Active Directory (Azure AD) and shared access signatures (SAS) only
	Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

Answer:

Blob storage:	Azure Active Directory (Azure AD) only
	Shared access signatures (SAS) only
	Access keys and shared access signatures (SAS) only
	Azure Active Directory (Azure AD) and shared access signatures (SAS) only
	Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)
File storage:	Azure Active Directory (Azure AD) only
	Shared access signatures (SAS) only
	Access keys and shared access signatures (SAS) only
	Azure Active Directory (Azure AD) and shared access signatures (SAS) only
	Azure Active Directory (Azure AD), access keys, and shared access signatures (SAS)

Explanation:

Blob storage:

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

File storage:

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

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: 142

☐☐ ☐☐☐ Azure ☐☐ ☐☐☐ ☐☐ ☐☐☐☐.

Name	Azure region
VM1	West Europe
VM2	West Europe
VM3	North Europe
VM4	North Europe

VM1☐ VM2☐ ☐☐☐☐ ☐☐ ☐☐☐ ☐☐☐ ☐☐☐☐.

Recovery Services☐ ☐☐☐☐ VM3 ☐ VM4☐ ☐☐☐☐ ☐☐☐.

☐☐ ☐☐ ☐☐☐ ☐☐ ☐☐☐?

A. ☐☐☐ ☐☐ ☐☐☐ ☐☐☐☐ ☐☐☐☐.

B. VM3 ☐ VM4☐ ☐☐ ☐☐☐ ☐☐☐☐☐.

C. ☐☐☐ ☐☐☐ ☐☐☐☐.

D. ☐☐☐ ☐☐ ☐☐☐ ☐☐☐☐.

Answer: B (LEAVE A REPLY)

Azure Backup uses the Recovery Services vault to protect virtual machines by creating restore points stored in Azure. Before a virtual machine can be backed up, it must have the Azure Backup extension (also called VMBackup extension) properly installed and configured on the VM.

According to Microsoft Learn - "Back up and restore encrypted virtual machines" and "Back up Azure VMs with the Azure Backup service", the backup process for a VM requires the following: "When you enable backup for an Azure virtual machine, the Azure Backup service installs an extension on the VM to coordinate the snapshot and backup process. If the extension is missing or not installed properly, backups cannot be triggered for that VM." Explanation of the Scenario You already have a Recovery Services vault protecting VM1 and VM2.

You now need to add VM3 and VM4 to the same vault for protection.

Before they can be added, you must ensure that:

The VMs are running in the same region as the Recovery Services vault (because a vault can only protect resources within its region).

The VMBackup extension is installed and registered on each VM.

If the backup extension is missing or not configured correctly, Azure Backup cannot communicate with the VM to take snapshots or register it in the vault.

Once the extension is properly configured, you can then enable backup from the Recovery Services vault without needing to create a new vault, policy, or storage account - because all those resources already exist and can be reused.

Other Options Explained

A). Create a new Recovery Services vault: # Not necessary unless the new VMs are in a different region from the existing vault.

C). Create a storage account: # Azure Backup automatically manages backup storage within the vault; no separate storage account is required.

D). Create a new backup policy: # Not needed unless you want a different backup schedule or retention. You can use the existing policy.

Final Verified Answer:

B). Configure the extensions for VM3 and VM4.

Microsoft Learn Extract (Supporting Documentation):

"To enable Azure Backup for virtual machines, the Azure Backup extension must be installed. This extension coordinates the snapshot and backup process and is installed automatically when backup is enabled."

"If the extension is not installed, you must configure it before enabling protection." (Source: Microsoft Learn - Back up Azure VMs with the Azure Backup service, Azure Recovery Services vault overview)

NEW QUESTION: 143

☐☐ ☐☐☐ App Service ☐☐☐ ☐☐ ☐☐☐☐☐.

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

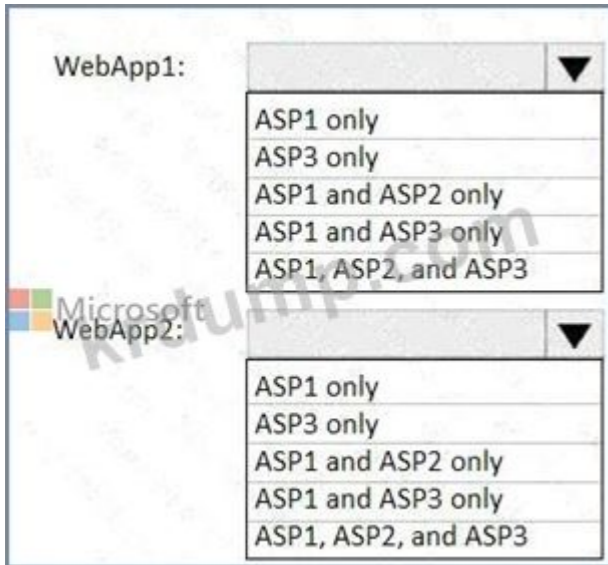
☐☐ ☐☐ ☐☐☐ Azure ☐☐☐ ☐☐ ☐☐☐☐☐.

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

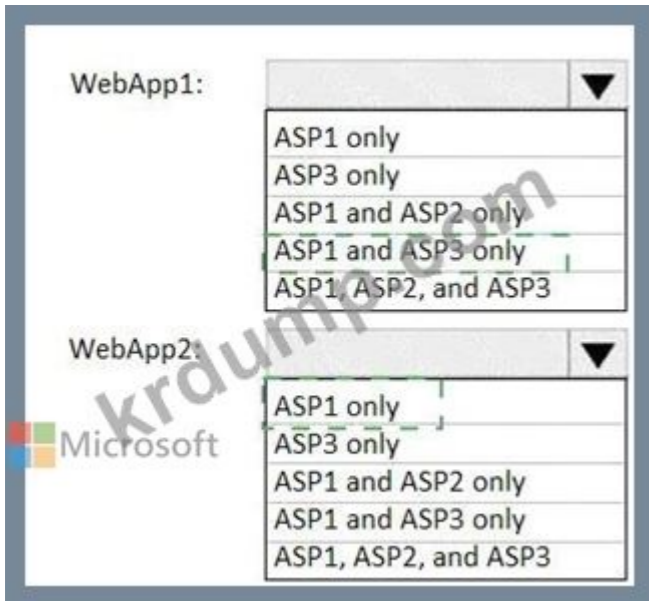
☐ ☐☐ ☐☐☐ ☐ ☐☐ App Service ☐☐☐ ☐☐☐☐ ☐☐☐.

☐☐☐ ☐☐☐☐ ☐☐☐? ☐☐☐☐☐☐ ☐☐ ☐☐☐☐☐ ☐☐☐ ☐☐☐ ☐☐☐☐☐☐.

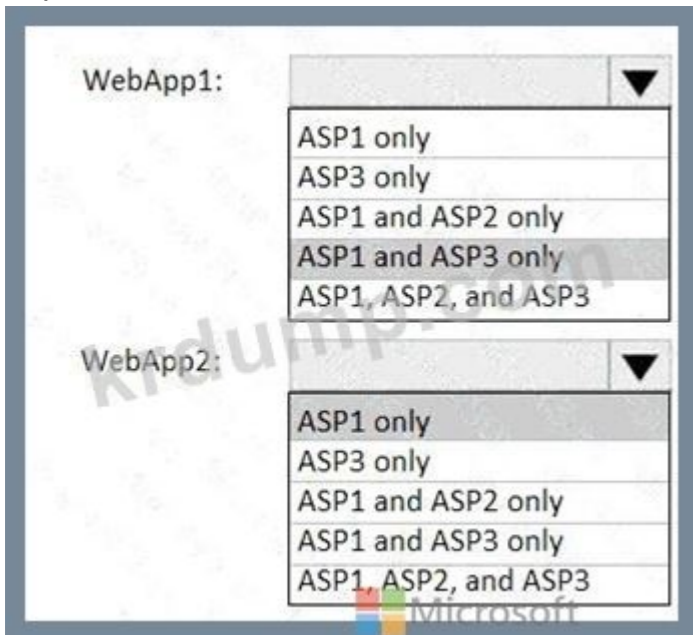
☐☐: ☐☐ ☐☐☐ 1☐☐☐☐.



Answer:



Explanation:



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

The App Service Plan and the Web App must be in the same region.

The operating system (Windows or Linux) of the App Service Plan must match the Web App type.

The runtime stack of the Web App must be supported on the App Service Plan's OS.

Let's apply these rules to the scenario:

App Service Plans:

Name

Operating System

Location

ASP1

Windows

West US

ASP2

Windows

Central US

ASP3

Linux

West US

Web Apps:

Name

Runtime Stack

Location

WebApp1

NET Core 3.0

West US

WebApp2

ASP.NET 4.7

West US

1## WebApp1 (.NET Core 3.0, West US)

Region requirement: Must use an App Service Plan in West US # # ASP1 and ASP3 qualify.

Runtime requirement: .NET Core 3.0 supports both Windows and Linux App Service Plans.

(Microsoft Learn: ".NET Core apps can run on both Windows and Linux App Service plans.")

Therefore, WebApp1 can use ASP1 (Windows, West US) or ASP3 (Linux, West US).

Answer: ASP1 and ASP3 only

2## WebApp2 (ASP.NET 4.7, West US)

Region requirement: Must be in West US # # ASP1 and ASP3 qualify.

Runtime requirement: ASP.NET 4.7 is a Windows-only framework; it cannot run on Linux App Service Plans.

(Microsoft Learn: "ASP.NET (non-Core) apps require a Windows-based App Service Plan.")

Therefore, only ASP1 (Windows, West US) is compatible.

Answer: ASP1 only

Final Verified Answers:

Web App

Compatible App Service Plans

WebApp1

ASP1 and ASP3 only

WebApp2

ASP1 only

Microsoft Documentation Extract (Azure App Service):

"App Service plans must be in the same region as the web app."

"Windows App Service plans host ASP.NET, .NET Core, and Node.js apps."

"Linux App Service plans host .NET Core, Node.js, Python, PHP, Java, and custom containers."

".NET Framework (ASP.NET 4.x) applications can only run on Windows-based App Service plans." Hence, the verified and Microsoft-official answer is:

WebApp1 # ASP1 and ASP3 only

WebApp2 # ASP1 only

NEW QUESTION: 144

adatum.com is an Azure DNS zone. research.adatum.com is a subdomain of adatum.com.

Which record type should be created in the adatum.com zone to delegate the research.adatum.com subdomain to another DNS server?

A. adatum.com PTR record for research.adatum.com.

B. adatum.com NS record for research.adatum.com.

C. adatum.com SOA record for research.adatum.com.

D. adatum.com A record for *.research.adatum.com.

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: 145

RG1 is a resource group in Azure. You have the following Bicep code in File1.bicep:

File1.bicep defines a virtual network (File1.bicep).

```
1 resource vnet 'Microsoft.Network/virtualNetworks@2023-11-01' = {
2   name: 'VNet1'
3   location: resourceGroup().location
4   tags: {
5     CostCenter: '12345'
6   }
7   properties: {
8     addressSpace: {
9       addressPrefixes: [
10        '10.0.0.0/24'
11      ]
12    }
13    enableVmProtection: false
14    enableDdosProtection: false
15    subnets: [
16      {
17        name: 'Subnet1'
18        properties: {
19          addressPrefix: '10.0.0.0/24'
20        }
21      }
22    ]
23  }
24 }
```

File2.bicep defines a virtual network (File2.bicep).

```
1 resource vnet 'Microsoft.Network/virtualNetworks@2023-11-01' = {
2   name: 'VNet1'
3   location: resourceGroup().location
4   tags: {
5     CostCenter: '67890'
6   }
7   properties: {
8     addressSpace: {
9       addressPrefixes: [
10        '10.0.0.0/16'
11      ]
12    }
13    enableVmProtection: false
14    enableDdosProtection: false
15    subnets: [
16      {
17        name: 'Subnet2'
18        properties: {
19          addressPrefix: '10.0.1.0/24'
20        }
21      }
22    ]
23  }
24 }
```

PowerShell

New-AzResourceGroupDeployment -ResourceGroupRG RG1 -TemplateFile File1.bicep New-AzResourceGroupDeployment -Whatif -ResourceGroupRG RG1 -TemplateFile File2.bicep

PowerShell commands for VNet configuration.

PowerShell commands for VNet configuration.

Answer Area

Statements	Yes	No
VNet1 has a CostCenter tag that has a value of 12345.	<input type="radio"/>	<input type="radio"/>
VNet1 has an IP address space of 10.0.0.0/16.	<input type="radio"/>	<input type="radio"/>
VNet1 has two subnets.	<input type="radio"/>	<input type="radio"/>

Answer:
ANSWER AREA

Statements	Yes	No
VNet1 has a CostCenter tag that has a value of 12345.	<input checked="" type="radio"/>	<input type="radio"/>
VNet1 has an IP address space of 10.0.0.0/16.	<input type="radio"/>	<input checked="" type="radio"/>
VNet1 has two subnets.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
VNet1 has a costcenter tag that has a value of 12345.	<input checked="" type="radio"/>	<input type="radio"/>
VNet1 has an IP address space of 10.0.0.0/16.	<input type="radio"/>	<input checked="" type="radio"/>
VNet1 has two subnets.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION: 146

Vault 1 Recovery Services Azure. Recovery Services vault configuration.

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

PowerShell commands for VM configuration.

Azure Backup configuration.

- A. VM1 □ VM2□ □□
- B. VM1□
- C. VM1 □ VM3□ □□
- D. VM1, VM2, VM3 □ VM4

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 147

□□ □□ □□□ □□□□ □□□ Azure □□□ □□□□.

Name	Type
LB1	Load balancer
VM1	Virtual machine
VM2	Virtual machine

LB1□ □□ □□ □□ □□□□.

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

□□ □□ □□□ □□□□ □□□ □□□□ NAT □□□ □□ □□□□□.

□□ 3389□ □□□□ □□□□□ VM2□ □□ □□ □□□□ □□□□ □□□□□.

- A. □□□□□ IP □□
- B. □□ □□□
- C. □□ □□ □□
- D. □□□ □

Answer: ([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 to. 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: 148

storage1□□□ Azure Storage □□□ □□□□.

Azure □□□□ □□□□□□ □□□□ App1□ App2□□ Azure App Service □□ □□□□. □ □□ □□□ ID□ □□□□□□.

App1□ App2□ storage1□□ BLOB□ □□ □ □□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.

* □□□□ □□□ □□ □□□□□□□.

* App2□ □□ 30□ □□ storage1□□□□ □□ □ □□□ □□□□□□.

□ □□ □□ storage1□□ □□□ □□□□ □□□□ □□□□ □□□ □□□ □□□□ □□□ □□□ □□□□□□. □□: □□ □□□ 1□□□□□.

Answer Area

App1:
Access keys
Advanced security
Access control (IAM)
Shared access signatures (SAS)

App2:
Access keys
Advanced security
Access control (IAM)
Shared access signatures (SAS)

Answer:
Answer Area

App1:
Access keys
Advanced security
Access control (IAM)
Shared access signatures (SAS)

App2:
Access keys
Advanced security
Access control (IAM)
Shared access signatures (SAS)

Explanation:

Answer Area

App1:

App2:

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: 149

☐☐ ☐☐ ☐☐☐☐ ☐☐☐ Azure ☐☐☐ ☐☐☐☐.

Name	Type
VMRG	Resource group
VNet1	Virtual network
VNet2	Virtual network
VM5	Virtual machine connected to VNet1
VM6	Virtual machine connected to VNet2

Azure☐☐ adatum.com☐☐☐ ☐☐ DNS ☐☐☐ ☐☐☐, VNet2☐ ☐☐ ☐☐☐☐ ☐☐☐ ☐☐☐☐, ☐☐ ☐☐☐ ☐☐☐☐☐☐.

adatum.com ☐☐☐ ☐☐ ☐☐☐ ☐☐ ☐☐☐☐ ☐☐☐☐.

Subscription (change)
Subscription ID
fde29b-d56a-4f6c-8298-6c53cd0b720c

ig5 (change)
click here to add tags

Search record sets

NAME	TYPE	TTL	VALUE
@	SOA	3600	Email: azuredns-hostmaster.microsoft.com Host: internal.cloudapp.net Refresh: 3600 Retry: 300 Expire: 2419200 Minimum TTL: 300 Serial number: 1
vm1	A	3600	10.1.0.4
vm9	A	3600	10.1.0.12

VM5 can resolve VM9.adatum.com.
VM6 can resolve VM9.adatum.com.

Answer Area

Microsoft

Statements	Yes	No
The A record for VM5 will be registered automatically in the adatum.com zone.	<input type="radio"/>	<input type="radio"/>
VM5 can resolve VM9.adatum.com.	<input type="radio"/>	<input type="radio"/>
VM6 can resolve VM9.adatum.com.	<input type="radio"/>	<input type="radio"/>

Answer:
Answer Area

Microsoft

Statements	Yes	No
The A record for VM5 will be registered automatically in the adatum.com zone.	<input type="radio"/>	<input checked="" type="radio"/>
VM5 can resolve VM9.adatum.com.	<input type="radio"/>	<input checked="" type="radio"/>
VM6 can resolve VM9.adatum.com.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
The A record for VM5 will be registered automatically in the adatum.com zone.	<input type="radio"/>	<input checked="" type="radio"/>
VM5 can resolve VM9.adatum.com.	<input type="radio"/>	<input checked="" type="radio"/>
VM6 can resolve VM9.adatum.com.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION: 150

Azure storage1. storage1 blob.

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

6. storage1 blob.

Name	Location	Access tier
File1	container1	Hot
File2	container2	Hot

On June 6, File1 will be stored in the Cool access tier.
On June 7, File2 will be stored in the Cool access tier.
On June 16, File2 will be stored in the Archive access tier.

Answer Area

Statements	Yes	No
On June 6, File1 will be stored in the Cool access tier.	<input type="radio"/>	<input type="radio"/>
On June 7, File2 will be stored in the Cool access tier.	<input type="radio"/>	<input type="radio"/>
On June 16, File2 will be stored in the Archive access tier.	<input type="radio"/>	<input type="radio"/>

Answer:

Answer Area

Statements	Yes	No
On June 6, File1 will be stored in the Cool access tier.	<input type="radio"/>	<input type="radio"/>
On June 7, File2 will be stored in the Cool access tier.	<input type="radio"/>	<input type="radio"/>
On June 16, File2 will be stored in the Archive access tier.	<input type="radio"/>	<input type="radio"/>

Explanation:

Answer Area

Statements	Yes	No
On June 6, File1 will be stored in the Cool access tier.	<input type="radio"/>	<input checked="" type="radio"/>
On June 7, File2 will be stored in the Cool access tier.	<input type="radio"/>	<input checked="" type="radio"/>
On June 16, File2 will be stored in the Archive access tier.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION: 151

Cluster1 is an Azure Kubernetes Service (AKS) cluster. The following table shows the IP addresses assigned to the cluster components.

IP address	Assigned to
131.107.2.1	Load balancer front end
192.168.10.2	Kubernetes DNS service
172.17.7.1	Docker bridge address
10.0.10.11	Kubernetes cluster node

Cluster1 is behind a load balancer. The load balancer front end IP address is 131.107.2.1.

When any internet user will try to access the cluster which is behind a load balancer, traffic will first hit to load balancer front end IP. So in the DNS configuration you have to provide the IP address of the load balancer.

- A. 172.17.7.1
- B. 131.107.2.1
- C. 192.168.10.2
- D. 10.0.10.11

Answer: B (LEAVE A REPLY)

When any internet user will try to access the cluster which is behind a load balancer, traffic will first hit to load balancer front end IP. So in the DNS configuration you have to provide the IP address of the load balancer.

Reference:

<https://stackoverflow.com/questions/43660490/giving-a-dns-name-to-azure-load-balancer>

AZ-104-KR ☐☐ ☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop ☐☐ ☐☐☐☐ ☐☐☐ AZ-104-KR ☐☐! DumpTop ☐ ☐☐ **AZ-104-KR** ☐☐ ☐☐☐ ☐☐☐☐☐☐, DumpTop AZ-104-KR ☐☐ ☐☐☐ ☐☐☐☐☐☐ ☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐. ☐☐☐☐☐ ☐☐☐☐☐ ☐☐☐☐☐ ☐☐ DumpTop AZ-104-KR ☐☐☐☐☐☐☐. <https://www.dumptop.com/Microsoft/AZ-104-KR-dump.html> (454 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 152

☐☐ ☐☐ ☐☐☐☐ ☐☐☐☐☐ ☐☐☐ Azure ☐☐☐☐ ☐☐☐☐☐.

Name	Type	Resource group	Tag
RG6	Resource group	<i>Not applicable</i>	<i>None</i>
VNET1	Virtual network	RG6	Department: D1

☐☐ ☐☐ ☐☐☐☐ ☐☐ RG6☐ ☐☐☐☐☐☐☐.

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

RG6☐ ☐☐☐☐☐☐☐: RGroup: RG6.

VNET2☐☐ ☐☐ ☐☐☐☐☐☐☐ RG6☐☐ ☐☐☐☐☐☐.

VNET1☐ VNET2☐ ☐☐ ☐☐☐☐☐☐☐? ☐☐☐☐☐☐☐ ☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐.

☐☐: ☐☐☐☐☐☐☐☐☐☐☐.


VNET1:

- None
- Department: D1 only
- Department: D1, and RGroup: RG6 only
- Department: D1, and Label: Value1 only
- Department: D1, RGroup: RG6, and Label: Value1


VNET2:

- None
- RGroup: RG6 only
- Label: Value1 only
- RGroup: RG6 and Label: Value1

Answer:

VNET1:  Microsoft ▼

None
 Department: D1 only
 Department: D1, and RGroup: RG6 only
 Department: D1, and Label: Value1 only
 Department: D1, RGroup: RG6, and Label: Value1

VNET2:  ▼

None
 RGroup: RG6 only
 Label: Value1 only
 RGroup: RG6, and Label: Value1

Explanation:

Answer Area

VNET1: Department: D1, and Label: Value1 only ▼

VNET2: Label: Value1 only ▼

 Microsoft

<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: 153

☐☐ ☐☐☐☐☐ Azure ☐☐☐☐ ☐ ☐ ☐☐☐☐.



Tenant Root Group



MG1



Sub1



RG1



VM1

Microsoft

□□□ □□□ □□□□ □□□ □□□ □□□ □□□□□.

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

□□: □□ □□□ 1□□□□.

Locks:

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

Tags:

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

Answer:

Locks:

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

Tags:

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

Explanation:

Locks:	<table border="1"> <tr><td> </td><td>▼</td></tr> <tr><td colspan="2">RG1 and VM1 only</td></tr> <tr><td colspan="2">Sub1 and RG1 only</td></tr> <tr><td colspan="2">Sub1, RG1, and VM1 only</td></tr> <tr><td colspan="2">MG1, Sub1, RG1, and VM1 only</td></tr> <tr><td colspan="2">Tenant Root Group, MG1, Sub1, RG1, and VM1</td></tr> </table>		▼	RG1 and VM1 only		Sub1 and RG1 only		Sub1, RG1, and VM1 only		MG1, Sub1, RG1, and VM1 only		Tenant Root Group, MG1, Sub1, RG1, and VM1	
	▼												
RG1 and VM1 only													
Sub1 and RG1 only													
Sub1, RG1, and VM1 only													
MG1, Sub1, RG1, and VM1 only													
Tenant Root Group, MG1, Sub1, RG1, and VM1													
Tags:	<table border="1"> <tr><td> </td><td>▼</td></tr> <tr><td colspan="2">RG1 and VM1 only</td></tr> <tr><td colspan="2">Sub1 and RG1 only</td></tr> <tr><td colspan="2">Sub1, RG1, and VM1 only</td></tr> <tr><td colspan="2">MG1, Sub1, RG1, and VM1 only</td></tr> <tr><td colspan="2">Tenant Root Group, MG1, Sub1, RG1, and VM1</td></tr> </table>		▼	RG1 and VM1 only		Sub1 and RG1 only		Sub1, RG1, and VM1 only		MG1, Sub1, RG1, and VM1 only		Tenant Root Group, MG1, Sub1, RG1, and VM1	
	▼												
RG1 and VM1 only													
Sub1 and RG1 only													
Sub1, RG1, and VM1 only													
MG1, Sub1, RG1, and VM1 only													
Tenant Root Group, MG1, Sub1, RG1, and VM1													

Box 1: Sub1, RG1, and VM1 only

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

Box 2: Sub1, RG1, and VM1 only

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

Reference:

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

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

NEW QUESTION: 154

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

Name	Location
VNet1	West Europe
VNet2	Southeast Asia
VNet3	South Central US

□□□□ □□ □□ □□□ □□□□ □□□□ □□□□.

Name	Virtual network	Service endpoint
Subnet1	VNet1	None
Subnet2	VNet2	Microsoft.Storage
Subnet3	VNet3	Microsoft.Storage
Subnet4	VNet3	None

□□□□ □□ □□ □□□ □□□ □□□ □□□□ □□□□.

Name	Location	Kind
storage1	West Europe	StorageV2
storage2	South Central US	BlobStorage
storage3	Southeast Asia	StorageV2

□□□ □□ □□□□ □□□ □□ □□□ □□□□ □□ □□ □□□ Azure □□□ policy1□□□ □□□ □□□□□ □□□ □□□□.

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

□□: □□ □□□ 1□□□□.

Answer Area

Statements	Yes	No
Policy1 can be applied to Subnet3.	<input type="radio"/>	<input type="radio"/>
Only storage1 and storage2 can be accessed from VNet2.	<input type="radio"/>	<input type="radio"/>
Only storage2 can be accessed from VNet3.	<input type="radio"/>	<input type="radio"/>

Answer:
ANSWER AREA

Statements	Yes	No
Policy1 can be applied to Subnet3.	<input checked="" type="radio"/>	<input type="radio"/>
Only storage1 and storage2 can be accessed from VNet2.	<input type="radio"/>	<input checked="" type="radio"/>
Only storage2 can be accessed from VNet3.	<input checked="" type="radio"/>	<input type="radio"/>

Explanation:

Policy1 can be applied to Subnet3. = YES

Only storage1 and storage2 can be accessed from VNet2. = NO

Only storage2 can be accessed from VNet3. = Yes

According to the Microsoft documentation, a service endpoint policy can be applied to any subnet in a virtual network that has a service endpoint enabled for the same service as the policy. In your scenario, Subnet3 has a service endpoint enabled for Microsoft.Storage, which is the same service as policy1. Therefore, policy1 can be applied to Subnet3.

According to the Microsoft documentation, when you configure network rules for a storage account, you can limit access to your storage account to requests that come from specified IP addresses, IP ranges, subnets in an Azure virtual network, or resource instances of some Azure services. In your scenario, storage1 and storage2 have network rules that allow access from Subnet1 and Subnet2 respectively. However, this does not mean that only these subnets can access the storage accounts. Other subnets or resources that have the same IP range or resource ID as Subnet1 or Subnet2

can also access the storage accounts. For example, Subnet4 in VNet2 has the same IP range as Subnet1 in VNet1, so it can also access storage1. Similarly, Subnet5 in VNet3 has the same IP range as Subnet2 in VNet1, so it can also access storage2. Therefore, only storage1 and storage2 cannot be accessed from VNet2.

According to the Microsoft documentation, when you create a private endpoint for a storage account, you assign a private IP address from your virtual network to the storage account. This enables secure traffic between your virtual network and the storage account over a private link. In your scenario, you have created a private endpoint for storage2 in Subnet6 of VNet3. This means that only Subnet6 can access storage2 over the private link. However, this does not mean that only Subnet6 can access storage2 at all. Other subnets or resources that have the same IP range or resource ID as Subnet6 can also access storage2 over the public endpoint of the storage account. For example, Subnet7 in VNet4 has the same IP range as Subnet6 in VNet3, so it can also access storage2 over the public endpoint. Therefore, only storage2 cannot be accessed from VNet3.

NEW QUESTION: 155

□□ □□ Azure □□□ □□ □□ □□ □□□ □□□ Azure □□□ □□□□.
□□ □□ □□□□ □□□□□□□ Azure Network Watcher □ □□□ □□□ □□□□ □□□.
□□ □ □□ □□□□ □□□□ □□□? □ □□□ □□□□ □□□ □□□□□.
□□: □□ □□□ 1□□□□□.

- A. Azure Monitor □ □□□ □□ □□(OCR)
- B. Log Analytics □□ □□
- C. Azure Monitor □□ □□
- D. □□ □□
- E. Microsoft Sentinel □□ □□

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: 156

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

□□□ □□□ □□□□□?

A. □

B. □□□

Answer: B (LEAVE A REPLY)

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

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

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

NEW QUESTION: 157

share1□□□ □□ □□□ □□□ storage1□□□ Azure Storage □□□ □□ □□□□□.

share!□ SMB □□□□□ □□□ □ □□□ □□□□ □□□. □□□□ □□□ □□□□□ □□□.

storage1□ □□□ □□□□ □□□?

A. □□ □□ □□□□(IRS)□ □□□ □□ □□

B. □□ □□ □□□□(LRS)□ □□ □□□□ □□

C. ZRS(Zone Redundant Storage)□ □□ □□ □□

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: 158

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

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

□□ □□ □□□ □□ □□□□ □□□□□.

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

☐☐☐ ☐ ☐☐ ☐☐☐ ☐☐☐☐ ☐☐ Vault1☐ ☐☐☐ ☐☐☐☐☐.

☐☐ ☐☐ ☐☐☐ Vault1☐ ☐☐☐ ☐ ☐☐☐?

- A. VM1, VM3, VMA ☐ VMC☐ ☐☐
- B. VM1 ☐ VM3☐ ☐☐
- C. VM1, VM2, VM3, VMA, VMB ☐ VMC
- D. VM1☐
- E. VM3 ☐ VMC☐ ☐☐

Answer: (SHOW ANSWER)

To create a vault to protect virtual machines, the vault must be in the same region as the virtual machines. If you have virtual machines in several regions, create a Recovery Services vault in each region.

References:

<https://docs.microsoft.com/bs-cyrl-ba/azure/backup/backup-create-rs-vault>

NEW QUESTION: 159

East US 2 ☐☐☐ VNET1☐☐☐ ☐☐ ☐☐☐☐☐☐ ☐☐☐ Azure ☐☐☐ ☐☐☐☐.

VM1-NI ☐ VM2-NI☐☐ ☐☐☐☐ ☐☐☐☐☐☐ VNET1☐ ☐☐☐☐☐.

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



```
Microsoft
{
  "apiVersion": "2024-07-01",
  "type": "Microsoft.Compute/virtualMachines",
  "name": "VM2",
  "zones": "2",
  "location": "EastUS2",
  "dependsOn": [
    "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
  ],
  "properties": {
    "computerName": "VM2",
    "adminUsername": "AzureAdmin",
    "adminPassword": "[parameters('adminPassword')]"
  },
  "storageProfile": {
```

```

"imageReference": "[variables('image')]",
"osDisk": {
  "createOption": "FromImage"
}
},
"networkProfile": {
  "networkInterfaces": [
    {
      "id": "[resourceId('Microsoft.Network/networkInterfaces', 'VM2-NI')]"
    }
  ]
}
}
}
}

```

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.

Statements	Yes	No
VM1 and VM2 can connect to VNET1.	<input type="radio"/>	<input type="radio"/>
If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.	<input type="radio"/>	<input type="radio"/>
If the East US 2 region becomes unavailable, VM1 or VM2 will be available.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
VM1 and VM2 can connect to VNET1.	<input checked="" type="radio"/>	<input type="radio"/>
If an Azure datacenter becomes unavailable, VM1 or VM2 will be available.	<input checked="" type="radio"/>	<input type="radio"/>
If the East US 2 region becomes unavailable, VM1 or VM2 will be available.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

East US.

NEW QUESTION: 161

Sub1 and Sub2 are Azure subscriptions, and Microsoft Entra ID is configured. VNet1, VNet2, VNet3, VNet4, and VNet5 are virtual networks.

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

Sub1 and Sub2 are Azure subscriptions, and Microsoft Entra ID is configured. VNet1, VNet2, VNet3, VNet4, and VNet5 are virtual networks.

VNet1 is connected to VNet2, VNet3, VNet4, and VNet5.

- A. VNet2 and VNet4 only
- B. VNet2, VNet3, and VNet4 only
- C. VNet2 only
- D. VNet2, VNet3, VNet4, and VNet5
- E. VNet2 and VNet3 only

Answer: D (LEAVE A REPLY)

NEW QUESTION: 162

cont2 is a container registry.

cont2 is configured with a storage policy. Which storage policy is used for the container images?

Options: 0, 1, 2, 3, 4, 5.



Answer:



krdump.com

Stored access policies:

1
2
3
4
5

Immutable blob storage policies:

1
2
3
4
5

Explanation:

Stored access policies: # 2

Immutable blob storage policies: # 1

From the scenario in the earlier case study (A. Datum Corporation):

- * The planned change requires creating a new container named cont2 in storage1, with the following:
- * Three stored access policies (Stored1, Stored2, Stored3)
- * A legal hold for immutable blob storage

Also, recall from the table:

- * storage1 has Hierarchical namespace = Yes (Data Lake Storage Gen2 enabled).
- * storage2 has Hierarchical namespace = No.

The question asks: "What is the maximum number of additional access policies you can create for cont2?" According to Microsoft Azure Storage documentation:

- * A blob container can have a maximum of five stored access policies.
- * These policies allow shared access signatures (SAS) to be managed centrally, letting you define start times, expiry times, and permissions at the container level.

Since three stored access policies (Stored1, Stored2, Stored3) already exist, the maximum additional policies that can still be created is:

5 (maximum allowed) - 3 (already existing) = 2

For immutable blob storage policies:

- * A container can have one active immutability policy at a time.
- * This can be either a time-based retention policy or a legal hold policy.
- * Since cont2 already has a legal hold applied, no additional immutable policy can coexist with it, but it can have one defined policy type (the legal hold).

Therefore:

- * Stored access policies: 2 additional policies can still be created.
- * Immutable blob storage policies: 1 policy (the existing legal hold).

This aligns exactly with Microsoft Learn: Azure Storage Blob Service limits and immutable storage documentation:

"A container may have up to five stored access policies. Immutable storage supports either a time-based retention policy or legal hold policy per container." Final Verified Answer:

Stored access policies: 2

Immutable blob storage policies: 1

NEW QUESTION: 163

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

VM1 D4s v3 500GB Puppet Agent
* 500GB
* Puppet Agent
* 500GB

VM1 500GB Puppet Agent

- A. 500GB
- B. D8s v3
- C. Puppet Agent
- D. 500GB

Answer: B (LEAVE A REPLY)

NEW QUESTION: 165

contoso.com Azure Active Directory(Azure AD) 500 CSV contoso.com Azure Portal Azure AD

- A.
- B.

Answer: B (LEAVE A REPLY)

In Microsoft Azure Active Directory (Azure AD), there is a clear distinction between internal users (members) and external users (guests). When you need to add a large number of external users (B2B collaborators) - such as those listed in a CSV file - you must use Azure AD B2B invitation processes, not the Bulk create user operation.

According to the Microsoft Azure Administrator Study Guide and Azure AD documentation, the Bulk create user operation in the Azure portal is designed only for internal user accounts within the organization's directory. It cannot be used to create guest user accounts. Guest users must be invited to the directory using either:

- * The New-AzureADMSInvitation PowerShell cmdlet, which sends invitations to external users' email addresses and creates guest accounts (userType = "Guest").
- * The Azure AD portal "Bulk invite" feature, which allows uploading a CSV file containing email addresses of external users to automate guest account creation.

The documentation explicitly states:

"To invite external users (B2B collaboration users) in bulk, you must use the Bulk invite feature or PowerShell with the New-AzureADMSInvitation cmdlet. The Bulk create operation is only supported for member users." (Source: Microsoft Learn - Azure Active Directory B2B collaboration and bulk operations guide) Therefore, since the scenario uses Bulk create user, it does not meet the goal of creating guest accounts for external users.

NEW QUESTION: 166

Adatum.com Microsoft Entra Subscription1 Azure Developers Subscription1 Dev Azure Logic Apps Subscription1 DevTest Labs

□□□ □□□ □□□□□?

A. □

B. □□□

Answer: B (LEAVE A REPLY)

The DevTest Labs User role is a built-in Azure role designed specifically to support operations within Azure DevTest Labs. According to the Microsoft Azure Administrator documentation on built-in roles, this role permits users to connect to, start, stop, restart, and use virtual machines inside a DevTest Lab, but it does not grant permissions to create or manage Azure resources outside of DevTest Labs, such as Azure Logic Apps.

In this scenario, the requirement is to allow the Developers group to create Azure Logic Apps in the Dev resource group. Azure Logic Apps are managed through the Microsoft.Logic resource provider and require permissions such as Microsoft.Logic/workflows/write. These permissions are included in roles like Logic App Contributor or the more general Contributor role at the resource group scope. The DevTest Labs User role does not include permissions to deploy or manage Logic Apps, even when assigned at the subscription level. Therefore, assigning this role does not meet the stated goal. Microsoft documentation explicitly states that DevTest Labs User is limited to DevTest Lab resources and cannot be used as a general deployment role.

Final Answer: B. No

AZ-104-KR □□ □□□ □□□□□ □□ DumpTop □□ □□□□ □□□ AZ-104-KR □□! DumpTop □ □□ **AZ-104-KR** □□ □□□ □□□□□□, DumpTop AZ-104-KR □□ □□□ □□□□□ □□□ □□□ □□□□□□□. □□□□ □□□ □□□□ □□ DumpTop AZ-104-KR □□□ □□□□□. <https://www.dumptop.com/Microsoft/AZ-104-KR-dump.html> (454 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)

NEW QUESTION: 167

storage1□□□ Azure Storage □□□ □□, □ □□□□ □□□□□ 1□ □□□□ 2□□ □ □□ □□□□□ □□□□. □ □□□□ □□□ Blob □□ □□□ □□□□□ □□□□. □□□□□ □□□ □□□ □□ □□□□ □□□□. □□□ □□ □□ □□ □□ □□□ □□□□.

regardless of the access tier of the base blob.

The rule for container2 has an action that moves blob versions to the Archive access tier if they are older than

30 days and have a prefix match of "archive/". Therefore, a blob version in container2 will only automatically move to the Archive access tier after 30 days if its name starts with "archive/". Otherwise, it will remain in its current access tier.

A rehydrated version is a blob version that was previously in the Archive access tier and was restored to an online access tier (Hot or Cool) by using the rehydrate priority option1. A rehydrated version does not automatically move to the Archive access tier after 30 days, unless there is a lifecycle management policy rule that explicitly specifies this action. In your case, neither of the rules applies to rehydrated versions, so they will stay in their online access tiers until you manually change them or delete them.

NEW QUESTION: 168

☐☐ ☐☐ ☐☐☐ ☐☐☐ ☐☐☐ ☐☐☐ Azure ☐☐☐ ☐☐☐☐.

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

Azure ☐☐☐☐ ☐☐☐ ☐☐☐☐☐☐☐ ☐☐☐☐ ☐☐ ☐☐☐☐ ☐☐☐ ZRS(☐☐ ☐☐ ☐☐☐☐) ☐☐☐ ☐☐☐ ☐☐☐☐ ☐☐☐☐☐☐.

☐☐☐ ☐☐☐☐ ☐☐☐?

- A. ☐☐☐☐1
- B. ☐☐☐2
- C. ☐☐☐3
- D. ☐☐4

Answer: B (LEAVE A REPLY)

<https://learn.microsoft.com/en-us/azure/storage/common/redundancy-migration?tabs=portal>

NEW QUESTION: 169

☐☐ ☐☐ ☐☐☐ ☐☐ ☐☐☐☐☐☐ ☐☐☐ Azure ☐☐☐ ☐☐☐☐.

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

☐☐☐☐ ☐☐ ☐☐ ☐☐☐ ☐☐ ☐☐☐ ☐☐☐☐☐☐☐☐.

Name	Connected to
VM1	VNet1
VM2	VNet2
VM3	VNet3

☐☐ ☐☐ ☐☐☐☐ ☐☐ IP ☐☐☐☐ ☐☐☐☐.

Bastion1☐☐☐ Azure Bastion ☐☐☐☐ VNet1☐☐☐☐☐☐.

Bastion1☐☐☐ ☐☐☐☐☐☐☐☐☐☐☐☐☐☐☐?

Answer: E (LEAVE A REPLY)

In Microsoft Azure, encryption scopes are a StorageV2 (general-purpose v2) storage account feature that allows fine-grained control over encryption settings for data stored within a single account. According to Microsoft Azure Storage documentation, an encryption scope defines a specific encryption context that can be applied at the container or blob level and is supported in non-hierarchical namespace storage accounts (those without Data Lake Gen2 enabled).

In the given scenario:

- * storage1 has Hierarchical namespace = Yes (Data Lake Storage Gen2 enabled).
- * storage2 has Hierarchical namespace = No.
- * The plan was to create an encryption scope named Scope1 in storage2.
- * The technical requirement specifies that Scope1 must be used to encrypt storage services.

According to the Azure Administrator documentation on encryption scopes:

"Encryption scopes are supported for block blobs, append blobs, page blobs, Azure Files, queues, and tables in standard StorageV2 accounts. Encryption scopes are not supported in hierarchical namespace (Data Lake Gen2) enabled accounts." This means that Scope1-created in storage2, which does not have hierarchical namespace-can encrypt all blob data (containers and blobs) as well as file shares, queues, and tables.

However, storage1 cannot use encryption scopes because hierarchical namespace storage accounts (ADLS Gen2) manage encryption at the account level and do not support per-scope encryption.

Therefore, only storage2 can apply Scope1, and it can encrypt containers, blobs, file shares, queues, and tables.

Topic 2, Contoso LtdOverview

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

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

Existing Environment

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

File servers

Domain controllers

Microsoft SQL Server servers

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

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

A SQL database

A web front end

A processing middle tier

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

Requirements

Planned Changes

Contoso plans to implement the following changes to the infrastructure:

Move all the tiers of App1 to Azure.

Move the existing product blueprint files to Azure Blob storage.

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

Technical Requirements

Contoso must meet the following technical requirements:

Move all the virtual machines for App1 to Azure.

Minimize the number of open ports between the App1 tiers.

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

Copy the blueprint files to Azure over the Internet.

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

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

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

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

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

Minimize administrative effort whenever possible.

User Requirements

Contoso identifies the following requirements for users:

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

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

Admin1 must receive email alerts regarding service outages.

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

NEW QUESTION: 172

Group4 is an Azure AD group. You need to grant Group4 read-only access to Azure file shares. Which Azure RBAC role should you assign to Group4?

A. storage1 or storage4

B. storage2

C. storage2 with ID-based access

D. storage1, storage2, storage4 with SAS

Answer: (SHOW ANSWER)

Azure role-based access control (Azure RBAC) for Azure file shares requires identity-based authentication integration. According to the Microsoft Azure Administrator documentation, this feature is only supported for StorageV2 (general purpose v2) and FileStorage account types.

In this scenario:

You are required to grant Group4 read-only access using Azure RBAC on Azure file shares.

The technical requirement specifies:

"Whenever possible, grant Group4 Azure RBAC read-only permissions to the Azure file shares." From the case study data:

Storage Account

Kind

Identity-based Access

storage1

Storage (general purpose v1)

Azure AD DS

storage2

StorageV2

Disabled

storage3

BlobStorage

N/A

storage4

FileStorage

Azure AD DS

The Storage (general purpose v1) type (storage1) does not support Azure AD or Azure RBAC integration for file shares. Microsoft documentation clearly states that "StorageV1 accounts must be upgraded to StorageV2 to support Azure AD authentication and RBAC role assignments." Meanwhile, FileStorage (storage4) already supports Azure AD Domain Services (Azure AD DS) and RBAC role assignment; hence no further modification is required there. However, to make storage1 compatible, it must be converted from StorageV1 to StorageV2.

Once converted to StorageV2, you can then:

Enable identity-based access for Azure file shares.

Assign Azure RBAC roles (e.g., Storage File Data Reader) to Group4.

Microsoft-Documented Requirements Summary:

Supported Account Types: StorageV2 or FileStorage

Unsupported: StorageV1 and BlobStorage

Required RBAC Roles for Read-Only Access:

Storage File Data Reader (or custom read-only role)

Thus, to meet the organization's requirement to provide Azure RBAC read-only permissions, you must change the account type of storage1 to StorageV2, ensuring both storage1 and storage4 can be managed with Azure RBAC.

NEW QUESTION: 173

Scenario: A company has a virtual network (VNET1) with two virtual machines (VM1 and VM2) and a load balancer (LB1). The load balancer is configured with two public IP addresses. The company wants to ensure that VM1 and VM2 can be added to LB1's backend pool.

Configuration details:

VM1 and VM2 are in the same VNet (VNET1).

LB1 is configured with two public IP addresses.

*LB1: LB1

*VM1: VM1

*SKU: Standard

*VNET1: VNET1

LB1 is configured with two public IP addresses.

VM1 and VM2 are in the same VNet (VNET1).

LB1 is configured with two public IP addresses?

A. No

B. Yes

Answer: B (LEAVE A REPLY)

This question tests understanding of Azure Load Balancer SKU compatibility and backend pool configuration requirements.

Scenario Summary

You have:

VM1 and VM2 in the same VNet (VNET1)

A Load Balancer (LB1) configured as:

Type: Internal

SKU: Standard

You need to ensure that VM1 and VM2 can be added to LB1's backend pool.

The proposed solution:

" You create two Standard public IP addresses and associate a Standard SKU public IP address to the network interface of each virtual machine. "

Understanding Azure Load Balancer Requirements

1. Backend pool requirements for a Standard Load Balancer:

All VMs must be in the same virtual network as the load balancer.

Each VM's NIC must be configured with a Standard SKU IP configuration (private or public).

The Load Balancer SKU must match the SKU of the IP addresses associated with the VM network interfaces.

2. Internal Load Balancer behavior:

An Internal Load Balancer (ILB) distributes traffic within a virtual network using private IP addresses, not public IPs.

Therefore, the backend VMs do not need public IPs - and adding them does not affect backend pool membership.

3. SKU alignment rule (Microsoft Docs):

"You can only attach virtual machines or instances that use Standard IP configurations to a Standard Load Balancer. Basic and Standard SKUs are not interchangeable." However:

A public IP is only required for inbound Internet access or outbound NAT, not for internal load balancing.

For an Internal Standard Load Balancer, backend pool members require Standard SKU NIC configurations, not public IPs.

Why the Proposed Solution Fails

The solution suggests creating two Standard public IPs and assigning them to the VMs' NICs.

This does not enable VM1 and VM2 to join the backend pool of an internal load balancer, because:

The load balancer type is internal, meaning it routes private traffic within the virtual network, not via public IPs.

Backend pool membership depends on the NIC's private IP configuration, not its public IP.

Adding public IPs only exposes VMs to the Internet and does not influence load balancer backend eligibility.

Thus, this action is unnecessary and does not meet the goal.

Correct Solution (for reference)

To meet the goal:

Ensure VM1 and VM2 have NICs configured with Standard SKU private IPs.

Ensure both VMs are in VNet1, the same virtual network as LB1.

No need to assign public IPs to internal backend VMs.

You could also ensure:

```
az network nic ip-config update \  
--name ipconfig1 \  
--nic-name VM1-nic \  
--resource-group RG1 \  
--private-ip-address-version IPv4 \  
--sku Standard
```

Final Verified Answer:

B). No

Microsoft Azure Documentation (Exam-Verified Extracts)

Azure Load Balancer SKU Comparison:

"Internal Load Balancer uses private IP addresses. Public IPs are not required or used for internal balancing." Backend Pool Membership:

"Virtual machines in the backend pool must be in the same virtual network as the load balancer and use matching Standard SKU IP configurations." Public vs Internal Load Balancer:

"For internal load balancers, only private frontends and backend configurations are supported."

Final Verified Answer: B. No

Assigning Standard public IPs to VMs does not affect internal load balancer backend connectivity. Backend membership depends on private IP configurations in the same VNet and matching SKU, not public IPs.

NEW QUESTION: 174

User1 is configuring storage1 in Azure. storage1 is a storage account. User1 is configuring storage1 with the following settings:

Name	Type
container1	Container
folder1	File share
Table1	Table

User1 is configuring storage1 with the following settings:

- * Blob
- * File
- * SMB

storage1 is a storage account. User1 is configuring storage1 with the following settings:

Allowed services 

Blob File Queue Table

Allowed resource types

Service Container Object

Allowed permissions

Read Write Delete List Add Create Update Process

Immutable storage

Blob versioning permissions

Enables deletion of versions

Allowed blob index permissions

Read/Write Filter

Start and expiry date/time

Start: 01/01/2022 12:00:00 PM

End: 01/01/2030 12:00:00 PM

(UTC+01:00) Belgrade, Bratislava, Budapest, Ljubljana, Prague

Allowed IP addresses

For example, 168.1.5.65 or 168.1.5.65-168.1.5.70

Allowed protocols

HTTPS only HTTPS and HTTP

Preferred routing tier

Basic (default) Microsoft network routing Internet routing

Some routing options are disabled because the endpoints are not published.

Signing key

key1

Generate SAS and connection string

User1 is configuring storage1 with the following settings:

key1:

- Table1 only
- Table1 and container1 only
- folder1 and Table1 only**
- folder1 and container1 only
- Table1, folder1, and container1

SAS1:

- Table1 only
- Table1 and container1 only**
- folder1 and Table1 only
- folder1 and container1 only
- Table1, folder1, and container1

Answer:

Answer Area  Microsoft

key1:

- Table1 only
- Table1 and container1 only
- folder1 and Table1 only**
- folder1 and container1 only
- Table1, folder1, and container1

SAS1:

- Table1 only
- Table1 and container1 only**
- folder1 and Table1 only
- folder1 and container1 only
- Table1, folder1, and container1

Explanation:

Answer Area

key1:

SAS1:

In this scenario, the Azure subscription contains a storage account (storage1) with three resources:

container1 (Blob container)

folder1 (File share)

Table1 (Table storage)

User1's assigned roles:

Storage Blob Data Reader # Read-only access to Blob data (cannot write).

Storage Table Data Contributor # Read, write, and delete access to Table data.

Storage File Data SMB Share Contributor # Read and write access to Azure File shares.

Now, let ' s analyze access using key1 and SAS1.

1## Access via key1

When using an account key, access is granted to all services within the storage account-Blob, File, Queue, and Table, because the key authenticates at the account level.

However, the question specifically asks:

"To which resources can User1 write by using SAS1 and key1?"

While account keys grant access to all resources, write permissions depend on the assigned roles of User1 for that account.

User1 can:

Write to the File share (folder1) because they are a Storage File Data SMB Share Contributor.

Write to the Table (Table1) because they are a Storage Table Data Contributor.

Cannot write to the Blob container (container1) because they only have the Storage Blob Data Reader role, which is read-only.

Therefore, using key1, User1 can write to folder1 and Table1 only.

2## Access via SAS1

The Shared Access Signature (SAS) shown in the exhibit specifies:

Allowed services: Blob, File, and Table (Queue not selected).

Allowed permissions: Read, Write, Delete, List, Add, Create, and Update.

Start and expiry: Valid from 2022 to 2030.

Protocols: HTTPS only.

SAS defines data plane access, allowing specific operations within the defined services, regardless of the user's assigned Azure roles.

Therefore, the SAS allows:

Access to Blob service # affects container1

Access to File service # affects folder1

Access to Table service # affects Table1

However, the question specifically distinguishes SAS1 and key1. Because SAS1 defines services explicitly, the question's options match how Azure limits SAS to only those selected services.

Hence, SAS1 allows write access to:

Table1 (Table service)

container1 (Blob service)

Therefore, using SAS1, User1 can write to Table1 and container1 only.

Final Verified Answer:

Access Method

Write Access To

key1

folder1 and Table1 only

SAS1

Table1 and container1 only

Reference Extract (Azure Documentation):

"A shared access signature (SAS) grants limited access rights to Azure Storage resources for a specified time and set of permissions. The permissions are defined per service selected (Blob, File, Queue, Table)."

"Using an account key provides full access to all data objects in the storage account, but effective access may still be limited by assigned RBAC roles if role-based access control is enforced."

NEW QUESTION: 175

Azure .

Azure Resource Manager(ARM) .

? .

: 1 .

Answer Area

```
{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "parameters": {
    "numberOfDataDisks": {
      "type": "int",
      "metadata": {
        "description": "The number of dataDisks to create."
      }
    }
  },
  ...
  "resources": [
    {
      "type": "Microsoft.Compute/virtualMachines",
      "apiVersion": "2017-03-30",
      ...
      "properties": {
        "storageProfile": {
          ...
          "copy": [
            {
              "copyIndex": {
                "dependsOn": [
                  "numberOfDataDisks"
                ]
            }
          ],
          ...
          "diskSizeGB": 1023,
          "lun": "[copyIndex] ('dataDisks')]",
          "createOptions": [
            "copyIndex",
            "copy",
            "dependsOn"
          ]
        }
      }
    }
  ]
}
```



Krcdump.com

Answer:

Answer Area



Microsoft

```
{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "parameters": {
    "numberOfDataDisks": {
      "type": "int",
      "metadata": {
        "description": "The number of dataDisks to create."
      }
    }
  },
  "resources": [
    {
      "type": "Microsoft.Compute/virtualMachines",
      "apiVersion": "2017-03-30",
      "properties": {
        "storageProfile": {
          "copy": [
            {
              "copyIndex": [copyIndex],
              "dependsOn": [
                {
                  "name": "[dataDisks]",
                  "type": "Microsoft.Compute/disks"
                }
              ]
            }
          ],
          "diskSizeGB": 1023,
          "lun": [copyIndex] ('dataDisks')],
          "createOptions": [copyIndex]
        }
      }
    }
  ]
}
```

Explanation:

```
Answer Area

{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "parameters": {
    "numberOfDataDisks": {
      "type": "int",
      "metadata": {
        "description": "The number of dataDisks to create."
      }
    }
  },
  ...
},
"resources": [
  {
    "type": "Microsoft.Compute/virtualMachines",
    "apiVersion": "2017-03-30",
    ...
    "properties": {
      "storageProfile": {
        ...
        "copy": {
          { "name": "dataDisks",
            "count": "[parameters('numberOfDataDisks')]",
            "input": {
              "diskSizeGB": 1023,
              "lun": "[copyIndex('dataDisks')]",
              "createOption": "Empty"
            }
          }
        }
      }
    }
  }
],
...
}
```

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: 176

VNet1 is a virtual network with a 10.0.0.0/16 IP address range. VNet1 has four subnets: Subnet0, Subnet1, Subnet2, and GatewaySubnet. Subnet0 has a 10.0.0.0/24 IP address range. Subnet1 has a 10.0.1.0/24 IP address range. Subnet2 has a 10.0.2.0/24 IP address range. GatewaySubnet has a 10.0.254.0/24 IP address range.

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 is connected to VM1. VM1 has a 10.0.1.10 IP address.

RT1 is a route table with a single route to Subnet1.

VM1 is connected to VNet1. VNet1 is connected to RT1.

RT1 has a route to Subnet1 with a metric of 1.

What is the IP address of VM1?

Answer Area Microsoft

Address prefix	10.0.0.0/16 10.0.1.0/24 10.0.254.0/24
Next hop type:	Virtual appliance Virtual network Virtual network gateway
Assigned to:	GatewaySubnet Subnet0 Subnet1 and Subnet2

Answer:

Answer Area Microsoft

Address prefix	10.0.0.0/16 10.0.1.0/24 10.0.254.0/24
Next hop type:	Virtual appliance Virtual network Virtual network gateway
Assigned to:	GatewaySubnet Subnet0 Subnet1 and Subnet2

Explanation:

Address prefix

▼
10.0.0.0/16
10.0.1.0/24
10.0.254.0/24

Next hop type:

▼
Virtual appliance
Virtual network
Virtual network gateway

Assigned to:

▼
GatewaySubnet
Subnet0
Subnet1 and Subnet2

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: 177

Microsoft 365 contoso.com Azure Active Directory(Azure AD) .

User1, User2, User3 Library1 Microsoft SharePoint .

. 180

.

.

A.

B. Office 365 . .

C. Office 365 . .

D.

E.

Answer: B,C (LEAVE A REPLY)

You can set expiration policy only for Office 365 groups in Azure Active Directory (Azure AD).

Note: With the increase in usage of Office 365 Groups, administrators and users need a way to clean up unused groups. Expiration policies can help remove inactive groups from the system and make things cleaner.

When a group expires, all of its associated services (the mailbox, Planner, SharePoint site, etc.) are also deleted.

You can set up a rule for dynamic membership on security groups or Office 365 groups.

NEW QUESTION: 178

Sub1 Azure .

.

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

.

*

* SQL

Explanation:

Answer Area

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

Protect the web servers from SQL injection attacks:

Microsoft

Box 1: an internal load balancer

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

Box 2: an application gateway that uses the WAF tier

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

NEW QUESTION: 179

□□ □□ □□□ □□□ □□□ □□□ Azure □□□ □□□□.

AzCopy□ □□□□ □□□ Blob□ □□ □□□ □□ □□□□□ □□□. AzCopy□ □□□ □ □□□ □□ □□□ □□□□ □□□.

□ □□□ □□ □□□ □□□□ □□□? □□□□ □□□ □□ □□□ □□ □□□□ □□□□□□□. □ □□□ □ □, □□ □ □□ □□ □□□□ □□ □ □□□□. □□□□ □□□ □ □□□ □□ □

□□ □□□□□□ □□□□□ □ □□ □□□□.

□□: □□ □□□ 1□□□□.

METHODS

- OAuth
- Anonymous
- A storage account access key
- A shared access signature (SAS) token

PREVIOUS PAGE

storage1:

storage2:

Microsoft

Answer:

Methods

- OAuth
- Anonymous
- A storage account access key
- A shared access signature (SAS) token

Answer Area

storage1: A shared access signature (SAS) token

storage2: A shared access signature (SAS) token

Microsoft

Explanation:

storage1: A shared access signature (SAS) token

storage2: A shared access signature (SAS) token

NEW QUESTION: 180

Q: You need to invite 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).
Q: How should you invite these users?
contoso.com Microsoft Entra ID portal.
500 users CSV file.
500 users contoso.com Microsoft Entra ID portal.
Q: Azure Portal Microsoft Entra ID portal?
Q: How should you invite these users?

- A. Microsoft Entra ID portal
- B. Microsoft Entra ID portal

Answer: A (LEAVE A REPLY)

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: 181

storage1 Blob Azure storage account.

contained Blob storage account.

Q: How should you configure storage1?

Q: How should you configure storage1?

- A. P-384 storage account EC

- B. P-521 EC
- C. P-384 EC
- D. 4096 RSA
- E. 2048, 3072, 4096 RSA

Answer: E (LEAVE A REPLY)

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

AZ-104-KR DumpTop **AZ-104-KR** DumpTop **AZ-104-KR**, DumpTop AZ-104-KR
<https://www.dumptop.com/Microsoft/AZ-104-KR-dump.html> (454 Q&As Dumps, **30%OFF**
Special Discount: KrDump)

NEW QUESTION: 182

Subscription1 Azure
 Subscription1 share1 Azure
 SAS1 (SAS)

Start

End

(UTC-06:00) Central Time (US & Canada)

Allowed IP addresses ⓘ

Allowed protocols ⓘ
 HTTPS only HTTPS and HTTP

Preferred routing tier ⓘ
 Basic (default) Microsoft network routing Internet routing

Some routing options are disabled because the endpoints are not published.

Signing key ⓘ

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



Answer Area

If on January 2, 2025, you run Microsoft Azure Storage Explorer on a computer that has an IP address of 193.77.134.1, and you use SAS1 to connect to the storage account, you [answer choice].

- will be prompted for credentials
- Will have no access
- will have read, write, and list access
- will have read-only access

If on January 10, 2025, you run the net use command on a computer that has an IP address of 193.77.134.50, and you use SAS1 as the password to connect to share1, you [answer choice].

- will be prompted for credentials
- will have no access
- Will have read, write, and list access
- will have read-only access

Answer:

Answer Area

If on January 2, 2025, you run Microsoft Azure Storage Explorer on a computer that has an IP address of 193.77.134.1, and you use SAS1 to connect to the storage account, you [answer choice].



Microsoft
krdump.com

will be prompted for credentials
 Will have no access
 will have read, write, and list access
 will have read-only access

If on January 10, 2025, you run the net use command on a computer that has an IP address of 193.77.134.50, and you use SAS1 as the password to connect to share1, you [answer choice].

will be prompted for credentials
 will have no access
 Will have read, write, and list access
 will have read-only access

Explanation:

Answer Area

If on January 2, 2025, you run Microsoft Azure Storage Explorer on a computer that has an IP address of 193.77.134.1, and you use SAS1 to connect to the storage account, you [answer choice].

will have no access

If on January 10, 2025, you run the net use command on a computer that has an IP address of 193.77.134.50, and you use SAS1 as the password to connect to share1, you [answer choice].

will have read, write, and list access

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

The exhibit shows the following important configuration details for SAS1:

Validity period: January 1, 2025 through January 1, 2028 # both access attempts occur within the valid timeframe.

Allowed IP addresses: Not specified # access is allowed from any IP address.

Allowed protocols: HTTPS only

Signing key: key1

Scenario 1 - Azure Storage Explorer

Azure Storage Explorer accesses Azure Storage using the HTTPS REST API. Because:

The access is within the SAS validity period,

No IP restriction is configured,

HTTPS is allowed,

the connection succeeds. However, the SAS configuration shown does not include write or list permissions, meaning access is limited to read-only operations.

Result: Will have read-only access

Scenario 2 - net use to connect to Azure file share

The net use command connects to Azure Files using SMB (port 445). SMB traffic does not use HTTPS.

Because the SAS explicitly restricts access to HTTPS only, SMB-based access is blocked, regardless of IP address or validity period.

Microsoft Azure documentation clearly states:

"When HTTPS only is selected, requests that use HTTP or SMB are denied."

Result: Will have no access

Answer Area



Minimum number of network interfaces:

5
10
15
20

Minimum number of network security groups:

1
2
5
10

Answer:

Answer Area

Minimum number of network interfaces:

5
10
15
20

Minimum number of network security groups:

1
2
5
10

Explanation:

Box 1: 5

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

Box 2: 1

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

Answer Area

Minimum number of network interfaces:

	▼
5	
10	
15	
20	

Minimum number of network security groups:

	▼
1	
2	
5	
10	

Reference:

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

NEW QUESTION: 185

App1 is a virtual machine (VM) in an Azure virtual network (VNet). App1 is connected to a virtual network interface card (NIC) on VM1. VM2 is also connected to the same NIC on VM1.

App1 is connected to a virtual network interface card (NIC) on VM1. VM2 is also connected to the same NIC on VM1.

VM2 is connected to a virtual network interface card (NIC) on VM1. VM1 is also connected to the same NIC on VM1.

Azure Backup VM1 backup schedule. Backup 01:00 every 30 minutes. What object should be used to configure the protection for VM1? Select one.

Answer Area

Location in which to store the backups:

- A blob container
- A file share
- A Recovery Services vault
- A storage account

Object to use to configure the protection for VM1:

- A backup policy
- A batch job
- A batch schedule
- A recovery plan

Answer:

Answer Area

Location in which to store the backups:

- A blob container
- A file share
- A Recovery Services vault
- A storage account

Object to use to configure the protection for VM1:

- A backup policy
- A batch job
- A batch schedule
- A recovery plan

Explanation:

Answer Area

Location in which to store the backups:

- A blob container
- A file share
- A Recovery Services vault
- A storage account

Object to use to configure the protection for VM1:

- A backup policy
- A batch job
- A batch schedule
- A recovery plan

Box 1: A Recovery Services vault

A Recovery Services vault is an entity that stores all the backups and recovery points you create over time.

Box 2: A backup policy

What happens when I change my backup policy?

When a new policy is applied, schedule and retention of the new policy is followed.

References:

<https://docs.microsoft.com/en-us/azure/backup/backup-configure-vault>

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-backup-faq>

A Recovery Services vault is a storage entity in Azure that houses data. The data is typically copies of data, or configuration information for virtual machines (VMs), workloads, servers, or workstations. You can use Recovery Services vaults to hold backup data for various Azure services such as IaaS VMs (Linux or Windows) and Azure SQL databases.

You can use backup policy to configure schedule.

<https://docs.microsoft.com/en-us/azure/backup/backup-azure-recovery-services-vault-overview>
<https://docs.microsoft.com/en-us/azure/backup/backup-azure-vm-first-look-arm>

NEW QUESTION: 187

Scenario: A company has an Azure subscription. The company has a virtual machine (VM) named VM1. The company wants to back up VM1. The company has a Recovery Services vault named Vault1. The company has a backup policy named Policy1. The company wants to configure the backup for VM1. The company wants to use Vault1 as the location in which to store the backups. The company wants to use Policy1 as the object to use to configure the protection for VM1.

Question: What should you do to configure the backup for VM1?

- A.
- B.

Answer: A (LEAVE A REPLY)

Your account must meet one of the following to enable traffic analytics:

Your account must have any one of the following Azure roles at the subscription scope: owner, contributor, reader, or network contributor.

Reference:

<https://docs.microsoft.com/en-us/azure/network-watcher/traffic-analytics-faq>

NEW QUESTION: 188

□□□□ □□□□ □□□□.

VNet1□ VNet2□□ □ □□ □□□□□ □□□ Azure □□□ □□□□. VNet1□ 192.168.8.0/24□ IP □□ □□□ □□□□, VNet2□ 192.168.9.0/24□ IP □□ □□□ □□□□□. □□ □□□□ □□□□ □□□. □□□□ □□ □□ □□□ □□□□ □□□.

* VNet1□ □□□□ VNet2□ □□□□ □□□ □ □□□ □□□□□.

* □-□□□□ □□□□□ □□□□ Azure □□□□ □□□ □ □□□ □□□□□.

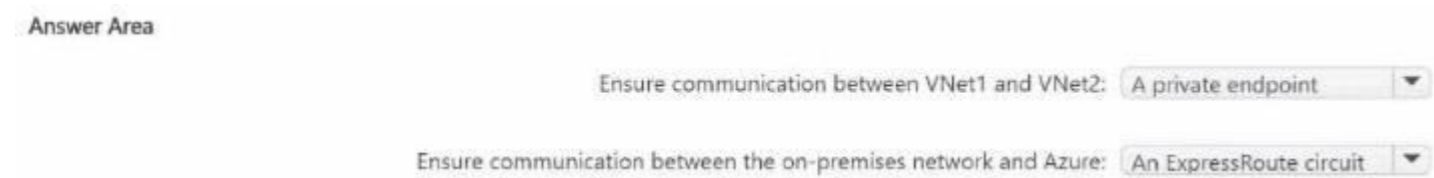
* □□□ □□□□□□□.



Answer:



Explanation:



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

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

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

NEW QUESTION: 189

□□ □□ □□□ □□□□ □□□ contoso.onmicrosoft.com□□□ Microsoft Entra □□□□ □□□□.

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

contoso.onmicrosoft.com (''').)

Self service password reset enabled

None Selected All

Select group
Group2

i These settings only apply to end users in your organization. Admins are always enabled for self-service password reset and are required to use two authentication methods to reset their password. Click here to learn more about administrator password policies.

.

(.)

Number of methods required to reset ⓘ

1 2

Methods available to users

- Mobile app notification
- Mobile app code
- Email
- Mobile phone
- Office phone
- Security questions

Number of questions required to register ⓘ

3 4 5

Number of questions required to reset ⓘ

3 4 5

Select security questions

10 security questions selected

i These settings only apply to end users in your organization. Admins are always enabled for self-service password reset and are required to use two authentication methods to reset their password. [Click here to learn more about administrator password policies.](#)



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

Answer Area

Statements

Yes

No

After User2 answers three security questions correctly, he can reset his password immediately.

If User1 forgets her password, she can reset the password by using the mobile phone app.

User3 can add security questions to the password reset process.



Answer:

Answer Area

Statements

Yes No

After User2 answers three security questions correctly, he can reset his password immediately.

If User1 forgets her password, she can reset the password by using the mobile phone app.

User3 can add security questions to the password reset process.

Explanation:

Answer Area

Statements

Yes No

After User2 answers three security questions correctly, he can reset his password immediately.

If User1 forgets her password, she can reset the password by using the mobile phone app.

User3 can add security questions to the password reset process.

NEW QUESTION: 190

□□ □□□ □□ Azure Storage □□□ □□□□.

Storage accounts

Contoso

+ Add Edit columns Refresh Assign Tags Delete

Subscriptions: All 2 selected - Don't see a subscription? Switch directories

Filter by name... All subscriptions All resource groups All types All locations No grouping

3 items

NAME	TYPE	KIND	RESOURCE	LOCATION	SUBSCRIPTI...	ACCESS T...	REPLICAT....
storageaccount1	Storage account	Storage	ContosoRG1	EastUS	Subscription 1	-	Read-access ge...
storageaccount2	Storage account	StorageV2	ContosoRG1	CentralUS	Subscription 1	Host	Geo-redundant...
storageaccount3	Storage account	BlobStorage	ContosoRG1	EastUS	Subscription 1	Host	Locally-redund....

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

□□: □□ □□□ 1□□□□.

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:
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

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.

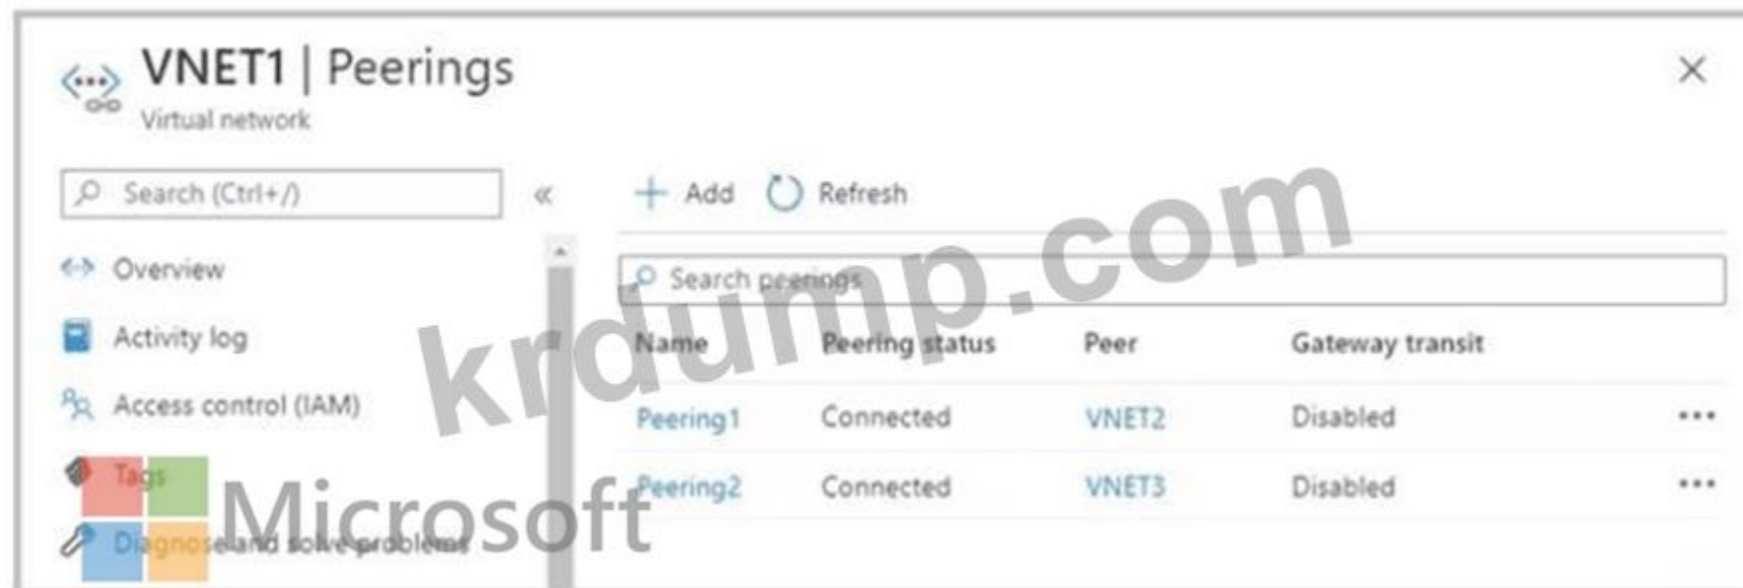
General-purpose v1 (GPv1) accounts provide access to all Azure Storage services, but may not have the latest features or the lowest per gigabyte pricing.

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

NEW QUESTION: 191

VNET1, VNET2, VNET3 are virtual networks in Azure.

VNET1 is peered with VNET2 and VNET3.



Answer Area



Packets from VNET1 can be routed to:

- VNET2 only
- VNET3 only
- VNET2 and VNET3

Packets from VNET2 can be routed to:

- VNET1 only
- VNET1 only
- VNET3 only
- VNET1 and VNET3

Answer:
 VNET1 and VNET3



Packets from VNET1 can be routed to:

- VNET2 only
- VNET3 only
- VNET2 and VNET3

Packets from VNET2 can be routed to:

- VNET1 only
- VNET1 only
- VNET3 only
- VNET1 and VNET3

Explanation:

Answer Area

Packets from VNET1 can be routed to: VNET2 and VNET3

Packets from VNET2 can be routed to: VNET1 only



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: 192

storage1 is a storage account in the Azure cloud. You need to ensure that storage1 can be accessed only by users who are members of the RBAC role 'Storage Blob Data Contributor'.

storage1 is a storage account in the Azure cloud. You need to ensure that storage1 can be accessed only by users who are members of the RBAC role 'Storage Blob Data Contributor'. Which of the following is the correct configuration for storage1?

A. Storage Blob Data Contributor

B. Storage Blob Data Reader

C. Storage Blob Data Owner

D. Storage Blob Data Contributor

E. Storage Blob Data Reader

F. Storage Blob Data Contributor

Answer: E (LEAVE A REPLY)

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

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

* Blob storage (containers and blobs)

* Queue storage

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

Conditions can restrict access to:

* Specific container names or blob prefixes.

* Specific queue names or messages.

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

Supported: Containers (Blob storage), Queues

Not supported: File shares, Tables

Microsoft Azure Reference (Conceptual Summary):

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

NEW QUESTION: 193

App1 and App2 are Azure App Service applications. App1 has a backup schedule of 1 day, starting on January 6, 2021, with a retention of 0 days. App2 has a backup schedule of 1 day, starting on January 6, 2021, with a retention of 30 days.

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

Which of the following statements are true?

Answer Area

Statements	Yes	No
On January 15, 2021, App1 will have only one backup in storage.	<input type="radio"/>	<input type="radio"/>
On February 6, 2021, you can access the backup of the App2 test slot from January 15, 2021.	<input type="radio"/>	<input type="radio"/>
On January 15, 2021, you can restore the App2 production slot backup from January 6 to	<input type="radio"/>	<input type="radio"/>



Answer:
answer area

Statements	Yes	No
On January 15, 2021, App1 will have only one backup in storage.	<input type="radio"/>	<input type="radio"/>
On February 6, 2021, you can access the backup of the App2 test slot from January 15, 2021.	<input type="radio"/>	<input type="radio"/>
On January 15, 2021, you can restore the App2 production slot backup from January 6 to	<input type="radio"/>	<input type="radio"/>



Explanation:
Answer Area

Statements	Yes	No
On January 15, 2021, App1 will have only one backup in storage.	<input checked="" type="radio"/>	<input type="radio"/>
On February 6, 2021, you can access the backup of the App2 test slot from January 15, 2021.	<input type="radio"/>	<input checked="" type="radio"/>
On January 15, 2021, you can restore the App2 production slot backup from January 6 to	<input checked="" type="radio"/>	<input type="radio"/>



On January 15, 2021, App1 will have only one backup in storage. Yes, this is correct. According to the table, App1 has a backup every 1 day, starting from January 6, 2021, with a retention of 0 days. This means that each backup will be deleted after 0 days, or as soon as the next backup is created. Therefore, on January 15, 2021, App1 will have only one backup in storage, which is the one created on that day1.

On February 6, 2021, you can access the backup of the App2 test slot from January 15, 2021. No, this is not correct. According to the table, App2 has a backup every 1 day, starting from January 6, 2021, with a retention of 30 days. This means that each backup will be deleted after 30 days, or when the storage limit is reached. However, the table also shows that App2 has a setting of "Keep at least one backup" set to Yes. This means that the oldest backup will be retained even if it exceeds the retention period or the storage limit2.

Therefore, on February 6, 2021, you can access the backup of the App2 test slot from January 6, 2021, but not from January 15, 2021.

On January 15, 2021, you can restore the App2 production slot backup from January 6 to the App2 test slot. Yes, this is correct. According to the web search results, you can restore a backup by overwriting an existing app or by restoring to a new app or slot3. You can also restore a backup from a different slot or app as long as they are in the same subscription and region4. Therefore, on January 15, 2021, you can restore the App2 production slot backup from January 6 to the App2 test slot.

NEW QUESTION: 194

VM1, VM2, VM3, and VM4 are in an Azure region. App1 is running on VM1.

App1 is configured with an Availability Zone. Which Availability Zone is App1 running in?

Which Availability Zone is App1 running in?

- A. Availability Zone 1
- B. Availability Zone 2
- C. Availability Zone 3
- D. Availability Zone 4

Answer: C (LEAVE A REPLY)

An Availability Zone in an Azure region is a combination of a fault domain and an update domain. For example, if you create three or more VMs across three zones in an Azure region, your VMs are effectively distributed across three fault domains and three update domains. The Azure platform recognizes this distribution across update domains to make sure that VMs in different zones are not updated at the same time.

Reference link

<https://learn.microsoft.com/en-us/training/modules/configure-virtual-machine-availability/5-review-availability-zones>

NEW QUESTION: 195

Which of the following is a valid time span for the Cost analysis blade?

- A. 7 days
- B. 30 days
- C. 90 days
- D. 180 days

Answer: C (LEAVE A REPLY)

Cost analysis: Correct Option

In cost analysis blade of Azure, you can see all the detail for custom time span. You can use this to determine expenditure of last few day, weeks, and month. Below options are available in Cost analysis blade for filtering information by time span: last 7 days, last 30 days, and custom date range. Choosing the first option (last 7 days) auditors can view the costs by time span.

Cost analysis shows data for the current month by default. Use the date selector to switch to common date ranges quickly. Examples include the last seven days, the last month, the current year, or a custom date range.

Pay-as-you-go subscriptions also include date ranges based on your billing period, which isn't bound to the calendar month, like the current billing period or last invoice. Use the <PREVIOUS and NEXT> links at the top of the menu to jump to the previous or next period, respectively. For example, <PREVIOUS will switch from the Last 7 days to 8-14 days ago or 15-21 days ago.



Invoice: Incorrect Option

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

Why don't I see an invoice for the last billing period?

There could be several reasons that you don't see an invoice:

- It's less than 30 days from the day you subscribed to Azure.
- The invoice isn't generated yet. Wait until the end of the billing period.
- You don't have permission to view invoices. If you have a Microsoft Customer Agreement, you must be the billing profile Owner, Contributor, Reader, or Invoice manager. For other subscriptions, you might not see old invoices if you aren't the Account Administrator. To learn more about getting access to billing information, see [Manage access to Azure billing using roles](#).
- If you have a Free Trial or a monthly credit amount with your subscription that you didn't exceed, you won't get an invoice unless you have a Microsoft Customer Agreement.

Resource Provider: Incorrect Option

When deploying resources, you frequently need to retrieve information about the resource providers and types. For example, if you want to store keys and secrets, you work with the Microsoft.KeyVault

resource provider. This resource provider offers a resource type called vaults for creating the key vault. This is not useful for reviewing all Azure costs from the past week which is required for audit.

Payment method: Incorrect Option

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

Reference:

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

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

NEW QUESTION: 196

Microsoft Entra ID Azure AD. Which of the following is true?

Name	On-premises sync enabled
User1	No
User2	Yes

Which of the following is true?

Name	Storage account
share1	contoso2024
share2	contoso2024
share3	contoso2025

contoso2024 ID is used to access share1.

The screenshot shows the 'File shares' configuration page in the Microsoft Entra ID console. The page title is 'CONTOSO2024 | ACTIVE DIRECTORY'. Below the title, there is a 'File shares' section with a 'Refresh' button. The main heading is 'Step 1: Enable an Active Directory source'. The instructions state: 'Choose the Active Directory source that contains the user accounts that will access a share in this storage account. You can set up identity-based access control for user accounts located in either one of these three domain services:'. A list of options is provided: 'Active Directory domain controller you host on a Windows Server (generally referred to as "on-premises AD" even though you might host these servers in Azure)', 'Azure Active Directory Domain Services (Azure AD DS), a platform as a service, hosted directory service and domain controller in Azure', and 'Azure AD Kerberos allows using Kerberos authentication from Azure AD-joined clients. In order to use Azure AD Kerberos, user accounts must be hybrid identities'. Below this list, there are three configuration cards: 'Active Directory' (Enabled, with a 'Configure' button), 'Azure Active Directory Domain Services' (Another access method is already configured), and 'Azure AD Kerberos' (Another access method is already configured). A blue information banner at the bottom states: 'Azure Active Directory (Azure AD) is not a domain controller, only a directory service. User accounts solely based in Azure AD are currently not supported.'

Step 2: Set share-level permissions

Once you have enabled Active Directory source on your storage account, you must configure share-level permissions in order to get access to your file shares. There are two ways you can assign share level permissions. You can assign them to all authenticated identities as a default share level permission and you can assign them to specific Azure AD users/user group. [Learn more](#)

Permissions for all authenticated users and groups

Default share-level permissions

Disable permissions and no access is allowed to file shares

Enable permissions for all authenticated users and groups

Select appropriate role *

Storage File Data SMB Share Contributor

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

Statements	Yes	No
User1 can access the content in share1.	<input type="radio"/>	<input type="radio"/>
User2 can access the content in share2.	<input type="radio"/>	<input type="radio"/>
User2 can access the content in share3.	<input type="radio"/>	<input type="radio"/>

Answer:

Statements	Yes	No
User1 can access the content in share1.	<input type="radio"/>	<input checked="" type="radio"/>
User2 can access the content in share2.	<input checked="" type="radio"/>	<input type="radio"/>
User2 can access the content in share3.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Statements	Yes	No
User1 can access the content in share1.	<input type="radio"/>	<input checked="" type="radio"/>
User2 can access the content in share2.	<input checked="" type="radio"/>	<input type="radio"/>
User2 can access the content in share3.	<input type="radio"/>	<input checked="" type="radio"/>

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

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

1. Azure Files Identity-Based Access Overview

Azure Files supports identity-based authentication and authorization through:

On-premises Active Directory Domain Services (AD DS)

Azure Active Directory Domain Services (Azure AD DS)

Azure AD Kerberos (hybrid identities required)

Important Note (Microsoft Docs):

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

2. Identity Sync (Hybrid Setup)

User

On-premises Sync Enabled

User1

No

User2

Yes

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

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

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

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

3. Storage Account Configuration

Share

Storage Account

share1

contoso2024

share2

contoso2024

share3

contoso2025

contoso2024

Configured with Active Directory (AD DS) integration (see exhibit).

Default share-level permissions are enabled for all authenticated users and groups with the Storage File Data SMB Share Contributor role.

This means any authenticated domain user (hybrid) has access.

contoso2025

No indication of AD DS configuration in the scenario.

Hence, it is not configured for identity-based access.

4. Step-by-Step Validation

User1 and share1 (contoso2024)

User1 is not hybrid (no on-prem sync).

SMB authentication requires Kerberos via domain-joined identity.

Result: # Cannot access share1.

User2 and share2 (contoso2024)

User2 is hybrid (on-prem sync enabled).

contoso2024 supports AD DS integration and allows authenticated domain users.

Result: # Can access share2.

User2 and share3 (contoso2025)

contoso2025 is not configured for AD DS integration.

□.

□□ □□ □□□ □□ □□□?

A. webapp1□□ □□□□□□ □□□ □□□□□.

B. webapp1□□ □□□ □□ □□□□ □□□□□.

C. plan1□□ App Service □□□ □□□□□.

D. plan1□□ App Service □□□ □□□□□.

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: 199

storageacct1234□□ □□□ □□□ □□□ User1, User2□□ □ □□ □□□□ □□□ Azure □□□ □□□□.

□□ □□□ □□□ □□□ User1□□ □□□□□.

The screenshot shows a window titled "User1 assignments - storageacct1234" with a close button (X) in the top right corner. Below the title is a subtitle: "Assignments for the selected user, group, service principal, or managed identity at this scope or inherited to this scope." There is a search bar with the placeholder text "Search by assignment name or description". Below the search bar, there are two sections: "Role assignments (2)" and "Deny assignments (0)". The "Role assignments (2)" section contains a table with the following data:

Role	Scope	Group assignment	Condition
Reader	Resource group (Inherited)	--	None
Storage Blob Data Contributor	This resource	--	Add

At the bottom of the window, there is a "Classic administrators (0)" section and the Microsoft logo.

User1□ □□□ □ □□ □ □□ □□□□□? □□□ □□ □□□ □□□□ □□□□□.

□□: □□ □□□ 1□□□□.

- A. storageacct1234□ □□ □□□ □□□□□.
- B. blob □□□□ storageacct1234□ □□□□□□.
- C. User2□ storageacct1234□ □□ □□□ □□□□□.
- D. storageacct1234□□ blob □□□□ □□□.
- E. storageacct1234□ □□□□ □□□□□.

Answer: (SHOW ANSWER)

In this scenario, User1 has two role assignments:

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

According to Microsoft Azure built-in roles documentation:

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

Authorization/* actions).

- * User1 cannot modify the firewall or network settings of the storage account - those actions require Contributor or Owner roles at the management plane level.
- * User1 can view file shares under the storage account metadata due to Reader access but cannot access data within them because that requires Storage File Data SMB Share Reader/Contributor permissions (different data plane).

Final Verified Answers:

- * B. Upload blob data to storageacct1234
- * D. View blob data in storageacct1234

Summary (from Microsoft Role Documentation Extracts):

- * Storage Blob Data Contributor = Can read/write/delete blob data (Data Plane)
- * Reader = Can view Azure resource configuration (Management Plane)
- * No permission to assign roles or modify security/firewall settings.

AZ-104-KR □□ □□□ □□□□□ □□ DumpTop □□ □□□□ □□□ AZ-104-KR □□! DumpTop □ □□ **AZ-104-KR** □□ □□□ □□□□□□, DumpTop AZ-104-KR □□ □□□ □□□□□ □□□ □□□ □□□□□□□. □□□□ □□□ □□□□ □□ DumpTop AZ-104-KR □□□ □□□□□. <https://www.dumptop.com/Microsoft/AZ-104-KR-dump.html> (454 Q&As Dumps, **30%OFF** Special Discount: **KrDump**)