



# Microsoft AZ-104 Exam Questions

**Total Questions: 530+**

**Demo Questions: 60**

**Version: Updated for 2025**

**Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner**

**For Access to the full set of Updated Questions – Visit:  
[AZ-104 Exam Dumps](#) by Cert Empire**

## Question: 1

### HOTSPOT

You have an Azure subscription that uses Azure Container Instances.

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

You create a container image named image1.

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

Which command should you run for each requirement? To answer, select the options in the answer area.

## Answer Area

Provision a new container registry:

▼

az acr build

az acr create

az container create

docker create

Add image1 to the registry:

▼

az acr create

az container create

docker pull

docker push

### Answer:

Provision a new Azure container registry: `az acr create --resource-group --name --sku`

Basic Add image1 to the registry: `docker push .azurecr.io/image1:latest`

### Explanation:

`az acr create` is the Azure CLI command specifically used to create (provision) an Azure Container Registry.

Once the local image is tagged for the registry and the user is logged in, `docker push .azurecr.io/` uploads (adds) the local image to that Azure Container Registry.

### References:

1. Microsoft Docs Azure CLI, `az acr create`: <https://learn.microsoft.com/cli/azure/acr#az-acrcreate>

(See Syntax: `az acr create` to create a new registry.)

<https://certempire.com>

2. Microsoft Docs Push an image to Azure Container Registry:

<https://learn.microsoft.com/azure/container-registry/container-registry-get-started-dockercli#push-image>

(Shows use of docker push .azurecr.io/imagename:tag.)

CertEmpire

## Question: 2

You have an Azure container registry named Registry1 that contains an image named image1. You receive an error message when you attempt to deploy a container instance by using image1. You need to be able to deploy a container instance by using image1.

Solution: You assign the AcrPull role to ACR-Tasks-Network for Registry1. Does this meet the goal?

A. Yes

B. No

### Answer:

B (No)

### Explanation:

The proposed solution is incorrect. While the AcrPull role provides the necessary permissions to pull an image from an Azure Container Registry (ACR), it must be assigned to the correct security principal. For an Azure Container Instance (ACI) to pull an image using role-based access control (RBAC), the AcrPull role should be assigned to the ACI's managed identity or a service principal that the ACI is configured to use. The principal ACR-Tasks-Network is related to the ACR Tasks feature for image building and is not the identity used by ACI for image deployment. Therefore, this assignment will not resolve the authentication error.

### References:

1. Microsoft Learn: Authenticate with Azure Container Registry from Azure Container Instances. This document explicitly states that to grant an ACI access to an ACR using RBAC, you must assign the AcrPull role to the ACI's managed identity or a service principal. It does not mention ACR-Tasks-Network as a valid principal for this purpose.

URL: <https://learn.microsoft.com/en-us/azure/container-instances/container-instances-authazure-container-registry#use-a-managed-identity>

2. Microsoft Learn: Azure built-in roles for Azure Container Registry. This page defines the AcrPull role, which grants "pull" permissions. It clarifies that this role is assigned to identities that need to pull images, such as container hosts.

URL: <https://learn.microsoft.com/en-us/azure/container-registry/container-registryroles?tabs=azur-e-cli#roles-and-permissions>

3. Microsoft Learn: About ACR Tasks. This documentation describes ACR Tasks as a suite of features for building, testing, and patching container images within Azure. This confirms that its related identities are for build processes, not for external services like ACI to pull

images for deployment.

URL: <https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tasksoverview>

CertEmpire

<https://certempire.com>

### Question: 3

You have an Azure container registry named Registry1 that contains an image named image1. You receive an error message when you attempt to deploy a container instance by using image1. You need to be able to deploy a container instance by using image1.

Solution: You select Use dedicated data endpoint for Registry1. Does this meet the goal?

A. Yes

B. No

#### Answer:

[No]

#### Explanation:

Enabling a dedicated data endpoint merely provides a separate FQDN for the registry's data-plane traffic. It neither affects authentication nor image-pull permissions. Azure Container Instances still must authenticate successfully (for example, via registry admin credentials or a managed identity) to pull image1. Therefore, the proposed action will not resolve the deployment error.

Why Incorrect Option is Wrong:

Yes Selecting a dedicated data endpoint does not change authentication or authorization; it therefore cannot fix a pull failure caused by credential or permission issues.

#### References:

1. Microsoft Docs Use dedicated data endpoints for Azure Container Registry, Overview, para. 2. <https://learn.microsoft.com/azure/container-registry/container-registry-data-endpoint>
2. Microsoft Docs Deploy container instances using images from Azure Container Registry, section Authenticate with a registry. <https://learn.microsoft.com/azure/containerinstances/container-instances-using-azure-container-registry>

## Question: 4

You have an Azure container registry named Registry1 that contains an image named image1. You receive an error message when you attempt to deploy a container instance by using image1. You need to be able to deploy a container instance by using image1.

Solution: You create a private endpoint connection for Registry1. Does this meet the goal?

A. Yes

B. No

### Answer:

B

### Explanation:

Creating a private endpoint for an Azure Container Registry (ACR) provides a private, secure network connection from a virtual network. However, it does not address the authentication requirement. To pull an image from a private registry, an Azure Container Instance (ACI) must be granted permission to access the registry. This is typically achieved by enabling the ACR's admin user and providing the credentials during ACI deployment, or by assigning a managed identity to the ACI with the AcrPull role on the registry. The proposed solution only solves for network access, not authentication, and therefore does not meet the goal on its own.

### References:

1. Azure Container Registry Authentication: Microsoft's official documentation outlines the necessary authentication methods for accessing a registry. It states, "To grant access to a registry, you assign permissions to an identity." Creating a private endpoint does not assign permissions.

Source: Microsoft Learn, "Authenticate with an Azure container registry," Section:

"Authentication options."

URL: <https://learn.microsoft.com/en-us/azure/container-registry/container-registryauthentication>

2. Deploying ACI from ACR: The documentation for deploying a container instance from a container registry explicitly details the need to provide credentials, either via the admin user account or a service principal/managed identity.

Source: Microsoft Learn, "Deploy to Azure Container Instances from Azure Container Registry," Section: "Deploy container with Azure CLI."

URL: <https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tutorialdeploy-ci>

3. ACR with Private Endpoints and ACI: When using a private endpoint, authentication is



still required. The documentation on this specific scenario confirms that network connectivity

via Private Link and authentication are separate, mandatory steps.

Source: Microsoft Learn, "Deploy a container instance into a virtual network," Section:  
"Deploy from Azure Container Registry."

URL: <https://learn.microsoft.com/en-us/azure/container-instances/container-instancesvnet#deploy-from-azure-container-registry>

CertEmpire

## Question: 5

You have a Standard Azure App Service plan named Plan1. You need to ensure that Plan1 will scale automatically when the CPU usage of the web app exceeds 80 percent. What should you select for Plan1?

- A. Automatic in the Scale out method settings
- B. Rules Based in the Scale out method settings
- C. Premium P1 in the Scale up (App Service plan) settings
- D. Standard S1 in the Scale up (App Service plan) settings
- E. Manual in the Scale out method settings

### Answer:

B

### Explanation:

To automatically scale an Azure App Service plan based on a performance metric like CPU usage, you must configure autoscale rules. This is achieved within the "Scale out (App Service plan)" settings by selecting the custom autoscale option, which allows you to define rules. A rule can be set to trigger a scale-out action (adding more instances) when the CPU percentage exceeds a specified threshold, such as 80%. This ensures the application has more resources to handle the increased load automatically.

### Why Incorrect Options are Wrong:

- A. Automatic in the Scale out method settings: While the goal is automatic scaling, the specific configuration method in Azure for metric-based scaling is defined by rules. "Rules Based" is the more precise term for this mechanism.
- C. Premium P1 in the Scale up (App Service plan) settings: This is a "scale up" operation, which changes the size and capabilities (CPU, memory) of the instances, not the number of instances. It does not provide automatic scaling based on load.
- D. Standard S1 in the Scale up (App Service plan) settings: This is a "scale up" setting. The plan is already Standard, and this action does not configure automatic scaling based on performance metrics.
- E. Manual in the Scale out method settings: This requires an administrator to manually change the instance count and is the opposite of the "automatic" scaling requirement.

### References:

Microsoft Azure Documentation: "Get started with autoscale in Azure". This document explains that custom autoscale operates based on rules that you define. It states, "You can <https://certempire.com>

configure autoscale rules that are based on metrics... For example, you can increase the number of instances when CPU usage is above 80%."

URL: <https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-get-started>

Microsoft Azure Documentation: "Scale instance count manually or automatically". This guide details the process for App Services, showing that to scale based on a metric, you must select "Custom autoscale" and "Add a rule".

URL: <https://learn.microsoft.com/en-us/azure/app-service/manage-scale-up#scale-instancecount-manually-or-automatically>

CertEmpire

## Question: 6

Overview -

ADatum Corporation is consulting firm that has a main office in Montreal and branch offices in Seattle and New York.

Existing Environment -

Azure Environment -

ADatum 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
Workspace1	Log Analytics workspace	RG1
WebApp1	Azure App Service web app	RG1
VNet1	Virtual network	RG2
zone1.com	Azure Private DNS zone	RG3

Azure Key Vault -

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

Vault1 contains the keys shown in the following table.

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

Microsoft Entra Environment -

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

Name	Microsoft Entra role	Azure role
Admin1	Global Administrator	<i>None</i>
Admin2	Attribute Definition Administrator	<i>None</i>
Admin3	Attribute Assignment Administrator	<i>None</i>
User1	<i>None</i>	Reader for RG2 and RG3

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 -

ADatum 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 Stored1, 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 Scope1.

Deploy new containers by using Image1 or Image2.

Technical Requirements -

ADatum 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.

You need to configure WebApp1 to meet the technical requirements.

Which certificate can you use from Vault1?

- A. Cert1 only
- B. Cert1 or Cert2 only
- C. Cert1 or Cert3 only
- D. Cert3 or Cert4 only
- E. Cert1, Cert2 Cert3, or Cert4

### Answer:

A

### Explanation:

CertEmpire

To configure a TLS binding for an Azure App Service using a certificate from Azure Key Vault, the certificate must satisfy two primary conditions:

1. It must be in the PFX (PKCS#12) file format, which is a container for the certificate and its corresponding private key.
2. The certificate's status within the Key Vault must be 'Enabled'.

Based on the provided table:

Cert1: Is in PFX format and is Enabled. It is a valid choice.

Cert2: Is in PEM format. This format is not supported for private key certificates used in App Service TLS bindings.

Cert3: Is Disabled. A disabled certificate cannot be accessed or used by Azure services.

Cert4: Is in PFX format and is Enabled. It is also technically a valid choice.

However, reviewing the available options, none of them list "Cert1 and Cert4". Options B, C, D, and E all include definitively invalid certificates (Cert2 or Cert3). By eliminating the incorrect options, Option A ("Cert1 only") is the only plausible choice, despite Cert4 also being technically valid.

### Why Incorrect Options are Wrong:

D: This option incorrectly includes the disabled Cert3 and omits the valid Cert1.

### References:

1. Microsoft Learn - Add a TLS/SSL certificate in Azure App Service: Under the section

<https://certempire.com>

"Import a certificate from Key Vault," the documentation states, "It must be a PFX certificate with a private key and uploaded to the vault." This confirms the PFX format requirement, invalidating Cert2.

URL: <https://learn.microsoft.com/en-us/azure/app-service/configure-ssl-certificate#import-a-certificate-from-key-vault>

2. Microsoft Learn - About Azure Key Vault certificates: This document describes the attributes of a Key Vault certificate object, including the enabled property. A certificate with enabled set to false is not operational. This invalidates Cert3.

URL: <https://learn.microsoft.com/en-us/azure/key-vault/certificates/aboutcertificates#certificate-attributes>

CertEmpire



## Question: 7

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

Solution: From the resource group blade, move VM1 to another subscription. Does this meet the goal?

A. Yes

B. No

### Answer:

B (No)

### Explanation:

Moving a virtual machine (VM) to a different subscription is an administrative action that changes the billing and management scope of the resource. This operation does not change the physical location (region) or the underlying host server where the VM is running. To force a VM to move to a new host to mitigate a maintenance event, you must perform an action that re-provisions the compute resources, such as redeploying the VM or stopping (deallocating) and then starting it again.

CertEmpire

### References:

1. Microsoft Azure Documentation - Move resources to a new resource group or subscription: "When you move a resource, it's moved to a new resource group or subscription. It doesn't change the location of the resource." This confirms that the physical location and host are unaffected by a subscription move.

URL: <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/moveresource-group-and-subscription>

2. Microsoft Azure Documentation - Redeploy virtual machine to new Azure node: "When you redeploy a VM, it moves the VM to a new node within the Azure infrastructure and then powers it back on... This task is useful if you are facing issues with remote desktop connection or application access to the VM." This is the correct procedure for moving a VM to a new host.

URL: <https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/redeploy-to-new-node-for-vm>

3. Microsoft Azure Documentation - Maintenance and updates for virtual machines in Azure: This document explains that for planned maintenance, Azure sometimes provides options for self-service maintenance. If immediate action is required, redeploying or

<https://certempire.com>

deallocating/restarting the VM are the standard user-initiated actions to move to a new host.

URL: <https://learn.microsoft.com/en-us/azure/virtual-machines/maintenance-and-updates>

CertEmpire

## Question: 8

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

Solution: From the VM1 Redeploy + reapply blade, you select Redeploy. Does this meet the goal?

A. Yes

B. No

### Answer:

A (Yes)

### Explanation:

The "Redeploy" action is the correct procedure for moving an Azure virtual machine to a new physical host within the Azure infrastructure. When a VM is redeployed, Azure shuts it down, migrates it to a different, healthy host node, and then powers it back on. This process is specifically designed to address issues with the underlying host, including planned maintenance, while retaining all configuration options and attached data disks. Therefore, using the "Redeploy" feature immediately moves VM1 to a different host, meeting the stated goal.

CertEmpire

### References:

1. Microsoft Learn. (2023). Redeploy Windows virtual machine to new Azure node. "When you redeploy a VM, it moves the VM to a new node within the Azure infrastructure and then powers it back on, retaining all your configuration options and associated resources."

<https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/redeploy-to-newnode-windows>

2. Microsoft Learn. (2023). Maintenance and updates for virtual machines in Azure. "If you want to proactively move your VM before the planned maintenance window, you can

redeploy your VM." <https://learn.microsoft.com/en-us/azure/virtual-machines/maintenanceand-updates#redeploy-your-vm>

## Question: 9

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

Solution: From the VM1 Updates blade, select One-time update. Does this meet the goal?

A. Yes

B. No

### Answer:

B (No)

### Explanation:

The "One-time update" feature, found within the "Updates" blade of a virtual machine, is a component of Azure Update Manager. Its purpose is to trigger an on-demand installation of available operating system (OS) updates within the guest VM. This action does not influence the underlying physical host allocation in the Azure infrastructure. To move a VM to a different host to preempt a maintenance event, an administrator must perform an action that re-provisions the VM on the Azure fabric, such as "Redeploy" or a "Stop (deallocate)" and "Start" cycle.

CertEmpire

### References:

1. Microsoft Learn - Redeploy virtual machine to new Azure node: This document explicitly states that the redeploy action moves the VM to a new node within the Azure infrastructure.

It is the correct procedure for this scenario.

URL: <https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/redeploy-tonew-node-windows>

2. Microsoft Learn - Planned maintenance for Azure virtual machines: This documentation explains that for maintenance events, users can sometimes proactively initiate the maintenance. The recommended action for moving a VM is to redeploy it.

URL: <https://learn.microsoft.com/en-us/azure/virtual-machines/maintenance-andupdates#self-service-maintenance>

3. Microsoft Learn - Deploy updates on a single VM: This document describes the "One-time update" feature, confirming it is used to "install updates on-demand on a single Azure VM," which pertains to OS patching, not host migration.

URL: <https://learn.microsoft.com/en-us/azure/update-center/deploy-updates-single-vm>

## Question: 10

Overview -

ADatum Corporation is consulting firm that has a main office in Montreal and branch offices in Seattle and New York.

Existing Environment -

Azure Environment -

ADatum 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
Workspace1	Log Analytics workspace	RG1
WebApp1	Azure App Service web app	RG1
VNet1	Virtual network	RG2
zone1.com	Azure Private DNS zone	RG3

Azure Key Vault -

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

Vault1 contains the keys shown in the following table.

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

Microsoft Entra Environment -

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

Name	Microsoft Entra role	Azure role
Admin1	Global Administrator	<i>None</i>
Admin2	Attribute Definition Administrator	<i>None</i>
Admin3	Attribute Assignment Administrator	<i>None</i>
User1	<i>None</i>	Reader for RG2 and RG3

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 -

ADatum 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 Stored1, 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 Scope1.

Deploy new containers by using Image1 or Image2.

Technical Requirements -

ADatum 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.

You need to meet the technical requirements for the KEK.

Which PowerShell cmdlet and key should you use?

- A. Set-AzVMDiskEncryptionExtension and Key2.
- B. Set-AzDiskEncryptionKey and Key2.
- C. Set-AzDiskDiskEncryptionKey and Key1.
- D. Set-AzVMDiskEncryptionExtension and Key1.

### Answer:

A

### Explanation:

The correct PowerShell cmdlet to enable Azure Disk Encryption (ADE) on a virtual machine is Set-AzVMDiskEncryptionExtension. This cmdlet includes parameters to specify a Key Encryption Key (KEK) from an Azure Key Vault for an added layer of security.

The technical requirements state to use a KEK whenever possible. For ADE, the KEK must be an RSA key. The provided table shows two valid keys in Vault1: Key1 (Type: RSA) and Key2 (Type: RSA-HSM). An RSA-HSM key provides superior security as the key material is safeguarded in a FIPS 140-2 validated hardware security module. Given that a more secure option (Key2) is available, it is the preferred choice.

### Why Incorrect Options are Wrong:

- B. Set-AzDiskEncryptionKey and Key2: Set-AzDiskEncryptionKey is not a valid PowerShell cmdlet for enabling Azure Disk Encryption on a virtual machine.
- C. Set-AzDiskDiskEncryptionKey and Key1: Set-AzDiskDiskEncryptionKey is not a valid PowerShell cmdlet; the name is syntactically incorrect and does not exist in the Az module.
- D. Set-AzVMDiskEncryptionExtension and Key1: While Key1 is a technically valid RSA key, Key2 (RSA-HSM) offers a higher level of security. Following security best practices, the HSM-backed key should be used when available.

### References:

1. Set-AzVMDiskEncryptionExtension Cmdlet: Microsoft. "Set-AzVMDiskEncryptionExtension." Azure PowerShell Documentation. Accessed May 20, <https://certempire.com>

2024. <https://learn.microsoft.com/en-us/powershell/module/az.compute/setazvmdiskencryptionextension>.

2. Azure Disk Encryption with KEK: Microsoft. "Azure Disk Encryption scenarios on Windows VMs - With a key encryption key (KEK)." Azure Virtual Machines Documentation.

Accessed May 20, 2024. <https://learn.microsoft.com/en-us/azure/virtualmachines/windows/disk-encryption-windows#with-a-key-encryption-key-kek>.

3. Key Vault Keys: Microsoft. "About Azure Key Vault keys, secrets and certificates." Azure Key Vault Documentation. Accessed May 20, 2024. <https://learn.microsoft.com/enus/azure/key-vault/general/about-keys-secrets-certificates#azure-key-vault-keys>. (This explains the difference between software and HSM-protected keys).

Topic: 5

CertEmpire



## Question: 11

### HOTSPOT -

You have an Azure subscription named Sub1.

You plan to deploy a multi-tiered application that will contain the tiers shown in the following table.

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

You need to recommend a networking solution to meet the following requirements:

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

Protect the web servers from SQL injection attacks.

Which Azure resource should you recommend for each requirement? To answer, select the appropriate options in the answer area.

### Answer Area

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

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

Protect the web servers from SQL injection attacks:

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

Hot Area:

### Answer:

1. Distribute traffic between web and business-logic VMs Internal Azure Load Balancer
2. Protect web tier from SQL-injection attacks Azure Application Gateway configured with Web Application Firewall (WAF)

### Explanation:

An internal (layer-4) Azure Load Balancer evenly load-balances TCP/UDP flows from the

web-server subnet to the middle-tier VM pool, satisfying the equal-spread requirement for

intra-VNet traffic.

SQL-injection prevention is provided by Azure Web Application Firewall, available on the WAF SKU of Azure Application Gateway; it inspects HTTP/S requests and blocks common OWASP threats such as SQLi before they reach the web servers.

**References:**

1. Microsoft Docs Azure Load Balancer distributes incoming network traffic across a pool of backend resources. (<https://learn.microsoft.com/azure/load-balancer/load-balanceroverview> , section What is Azure Load Balancer?)
2. Microsoft Docs Web Application Firewall (WAF) on Azure Application Gateway protects web applications from common threats such as SQL injection (<https://learn.microsoft.com/azure/web-application-firewall/ag/ag-overview> , first paragraph)

CertEmpire

## Question: 12

Your company has three offices. The offices are located in Miami, Los Angeles, and New York. Each office contains datacenter. You have an Azure subscription that contains resources in the East US and West US Azure regions. Each region contains a virtual network. The virtual networks are peered. You need to connect the datacenters to the subscription. The solution must minimize network latency between the datacenters. What should you create?

- A. three Azure Application Gateways and one On-premises data gateway
- B. three virtual hubs and one virtual WAN
- C. three virtual WANs and one virtual hub
- D. three On-premises data gateways and one Azure Application Gateway

### Answer:

B

### Explanation:

Azure Virtual WAN is a networking service designed for large-scale, optimized, and automated branch-to-branch and branch-to-Azure connectivity. The solution involves creating a single Virtual WAN to act as the management entity. Within this Virtual WAN, you deploy virtual hubs in Azure regions geographically close to your on-premises datacenters (e.g., East US for Miami/New York, West US for Los Angeles). These hubs are automatically interconnected in a full mesh over the Microsoft global backbone. This architecture enables low-latency, transitive routing between all connected datacenters, directly fulfilling the requirement to minimize network latency between the offices.

### Why Incorrect Options are Wrong:

- A: Azure Application Gateway is a Layer 7 load balancer for web traffic, and an On-premises data gateway is for service-level integration (e.g., Power BI), not for network-level site connectivity.
- C: This option inverts the correct architecture. A Virtual WAN is the top-level resource that contains one or more virtual hubs, not the other way around.
- D: This option uses the same incorrect components as option A. These services are not designed for establishing site-to-site network connections.

### References:

1. Microsoft Azure Documentation: "What is Azure Virtual WAN?". This document states, "Azure Virtual WAN is a networking service that brings many networking, security, and routing functionalities together... These functionalities include branch connectivity (via <https://certempire.com>

connectivity automation from Virtual WAN Partner devices such as SD-WAN or VPN CPE), Site-to-site VPN connectivity, remote user VPN (Point-to-site) connectivity, private (ExpressRoute) connectivity, intra-cloud connectivity (transitive connectivity for virtual networks), VPN ExpressRoute inter-connectivity, and routing..."

URL: <https://learn.microsoft.com/en-us/azure/virtual-wan/virtual-wan-about>

2. Microsoft Azure Documentation: "Virtual WAN architecture". This page illustrates the architecture, showing a single Virtual WAN resource that contains multiple hubs in different regions. It explicitly describes the "any-to-any connectivity" capability, which is key to minimizing latency between the datacenters.

URL:

<https://learn.microsoft.com/en-us/azure/virtual-wan/migrate-to-virtual-wan#keyconsiderations> (See "Hub-and-spoke vs. Virtual WAN" section)

3. Microsoft Azure Documentation: "Tutorial: Create a site-to-site connection using Azure Virtual WAN". This tutorial demonstrates the correct creation process: first, create the Virtual WAN, and then create a hub within it. This confirms the one-to-many relationship between a Virtual WAN and its hubs.

URL: <https://learn.microsoft.com/en-us/azure/virtual-wan/virtual-wan-site-to-site-portal> (See Step 1: "Create a virtual WAN" and Step 2: "Create a hub")

## Question: 13

### HOTSPOT -

You plan to deploy five virtual machines to a virtual network subnet.

Each virtual machine will have a public IP address and a private IP address.

Each virtual machine requires the same inbound and outbound security rules.

What is the minimum number of network interfaces and network security groups that you require? To answer, select the appropriate options in the answer area.

## Answer Area

Minimum number of network interfaces:

	▼
5	
10	
15	
20	

Minimum number of network security groups:

	▼
1	
2	
5	
10	

Hot Area:

### Answer:

Minimum number of network interfaces: 5 Minimum number of network security groups: 1

### Explanation:

Each Azure virtual machine requires its own network interface (NIC) to connect to a virtual network and receive a private IP address. To have a public IP, it must be associated with the NIC. Therefore, five VMs require a minimum of five NICs.

A single Network Security Group (NSG) contains a set of security rules that can be applied to multiple resources. Since all five VMs require the same security rules, one NSG can be created and associated with the subnet where all the VMs reside. This single NSG will enforce the common rules for all five VMs, making it the most efficient configuration.

### References:

<https://certempire.com>

1. Microsoft Azure Documentation, Create, change, or delete a network interface: "An Azure virtual machine (VM) must have at least one network interface (NIC) attached to it. A VM can have more than one NIC, depending on the size of the VM you create." This confirms the one-to-one minimum relationship between a VM and a NIC.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-networkinterface>

2. Microsoft Azure Documentation, Network security groups: "You can associate one network security group to zero, or one, or more network interfaces and subnets. [...] For simpler management of security rules, it's recommended that you associate an NSG to a subnet rather than individual NICs within the subnet." This confirms that a single NSG can be used for multiple resources.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groupsoverview>

CertEmpire

## Question: 14

You have an Azure subscription that contains the resources shown in the following table. LB1 is configured as shown in the following table.

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

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

Provide Remote Desktop access to VM1 from the internet by using port 3389.

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

What should you create on LB1 before you can create the new inbound NAT rules?

CertEmpire

- A. a frontend IP address
- B. a load balancing rule
- C. a health probe
- D. a backend pool

**Answer:**

A

**Explanation:**

An inbound NAT rule on an Azure Load Balancer maps traffic from a specific frontend IP address and port to a specific virtual machine in the backend pool. Each rule must have a unique combination of frontend IP and frontend port.

The requirement is to provide Remote Desktop access (port 3389) to both VM1 and VM2. To create two separate inbound NAT rules that both use the same frontend port (3389), you must use two different frontend IP addresses. Since the load balancer, LB1, currently has only one frontend IP address (pip1), you must create a second frontend IP configuration before you can create the two required inbound NAT rules.

**Why Incorrect Options are Wrong:**

<https://certempire.com>



B. a load balancing rule: A load balancing rule is for distributing traffic across a pool of VMs,

not for directing traffic to a specific VM, which is the function of an inbound NAT rule.

C. a health probe: Health probes are used by load balancing rules to determine the health of backend instances. They are not a prerequisite for creating inbound NAT rules.

D. a backend pool: A backend pool (pool1) that contains both target virtual machines (VM1 and VM2) already exists. Therefore, a new one is not required.

## References:

Microsoft Azure Documentation (Load Balancer Components): "An inbound NAT rule forwards incoming traffic sent to a frontend IP address and port combination. The traffic is sent to a specific virtual machine or instance in the backend pool." This highlights the mapping to a specific VM.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/components#inbound-nat-rule>

Microsoft Azure Documentation (Manage Inbound NAT Rules): The portal and CLI examples for creating inbound NAT rules show that you must specify a unique frontend port for each rule associated with a single frontend IP configuration. To reuse a port, a different frontend IP is necessary.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/manage-inbound-nat-rules>

Microsoft Azure Documentation (Frontend IP Configuration): "The frontend IP address of a load balancer... You can have multiple frontend IPs. Each frontend IP must have a public IP address." This confirms that multiple frontends can be added to a load balancer.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/components#frontend-ipconfiguration>

## Question: 15

### HOTSPOT -

You have Azure virtual machines that run Windows Server 2019 and are configured as shown in the following table.

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

You create a private Azure DNS zone named adatum.com. You configure the adatum.com zone to allow auto registration from VNET1.

Which A records will be added to the adatum.com zone for each virtual machine? To answer, select the appropriate options in the answer area.

### Answer Area

A records for VM1:

▼

None

Private IP address only

Public IP address only

Private IP address and public IP address

A records for VM2:

▼

None

Private IP address only

Public IP address only

Private IP address and public IP address

Hot Area:

### Answer:

A records for VM1: Private IP address only

A records for VM2: None

### Explanation:

Azure DNS private zone auto-registration for a virtual machine is contingent on two primary conditions:

1. The virtual network containing the VM must be linked to the private DNS zone for registration.
2. The VM's primary DNS suffix, as configured within its operating system, must

match the name of the private DNS zone.

In this scenario, VNET1 is linked to the adatum.com private zone.

- 

VM1: The DNS suffix is Adatum.com, which matches the zone name.

Therefore, it will automatically register an A record for its private IP address (10.1.0.4) in the adatum.com zone. Azure private DNS zones only manage records for private IP addresses, not public ones.

- 

VM2: Although it is in the linked VNET1, its DNS suffix is Contoso.com. Since this does not match the zone name adatum.com, no A record will be automatically created for VM2 in this zone.

## References:

- 

Microsoft Azure Documentation | What is an Azure private DNS zone?:

This official documentation explains the functionality of Azure Private DNS.

Under the section "Automatic registration," it states, "When you link a virtual network to a private DNS zone, you can enable automatic registration of DNS records for the virtual machines in that virtual network... If you enable automatic registration, Azure DNS will create an A record in the private zone for each virtual machine in the linked virtual network. The A record will map the virtual machine's name to its private IP address." This confirms that only the private IP is registered.

o Direct URL: <https://docs.microsoft.com/en-us/azure/dns/private-dns-overview>

- 

Microsoft Azure Documentation | DNS settings on virtual machines: This document clarifies the requirement for the DNS suffix. It notes that for a VM to register with a private DNS zone, its connection-specific DNS suffix must match the private DNS zone name.

o Direct URL: [https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-name-resolution-for-vms-and-role-](https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-name-resolution-for-vms-and-role-instances#dns-settings-on-virtual-machines)

[instances#dns-settings-on-virtual-machines](https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-name-resolution-for-vms-and-role-instances#dns-settings-on-virtual-machines) (Refer to the section on DNS suffix configuration).

## Question: 16

### HOTSPOT -

You have an Azure virtual network named VNet1 that connects to your on-premises network by using a site-to-site VPN. VNet1 contains one subnet named Subnet1.

Subnet1 is associated to a network security group (NSG) named NSG1. Subnet1 contains a basic internal load balancer named ILB1. ILB1 has three Azure virtual machines in the backend pool.

You need to collect data about the IP addresses that connects to ILB1. You must be able to run interactive queries from the Azure portal against the collected data.

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

### Answer Area

Resource to create:

	▼
An Azure Event Grid	
An Azure Log Analytics workspace	
An Azure Storage account	

Resource on which to enable diagnostics:

	▼
ILB1	
NSG1	
The Azure virtual machines	

Hot Area:

### Answer:

To collect the data, you should create a: Log Analytics workspace

To configure the data collection, you should configure: the NSG flow logs for NSG1

### Explanation:

To meet the requirement of running interactive queries against collected data, a Log Analytics workspace is necessary. It serves as the central repository for Azure Monitor Logs and enables analysis using the Kusto Query Language (KQL) directly from the Azure portal. To collect data about IP addresses connecting to the internal load balancer (ILB1), you must capture the network traffic flow. NSG flow logs are the specific feature designed for this purpose. By enabling flow logs on the Network Security Group (NSG1) associated with the load balancer's subnet (Subnet1), you can record all ingress and egress IP traffic. These logs can then be sent to the Log Analytics workspace for storage and interactive querying.

## References:

Log Analytics workspace: Microsoft Learn. (2023). Log Analytics workspace overview. "A Log Analytics workspace is a unique environment for Azure Monitor log data... The workspace provides a geographic location for the data, data isolation, and scope for configurations like data retention... You analyze this data by running log queries in Log Analytics..."

URL: <https://learn.microsoft.com/en-us/azure/azure-monitor/logs/log-analytics-workspaceoverview>

NSG Flow Logs: Microsoft Learn. (2023). Introduction to flow logging for network security groups. "Network security group (NSG) flow logging is a feature of Azure Network Watcher that allows you to log information about IP traffic flowing through an NSG... Flow log data is sent to Azure Storage accounts. From there, you can access the data and export it to any visualization tool, SIEM, or IDS of your choice."

URL: <https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-nsg-flowlogging-overview>

NSG Flow Logs with Log Analytics: Microsoft Learn. (2024). Tutorial: Log network traffic to and from a virtual machine using the Azure portal. "You can use network security group (NSG) flow logs to log network traffic... You can also analyze the flow logs by using Traffic Analytics, which aggregates flow log data and combines it with security information to provide a comprehensive view of your network traffic." Traffic Analytics requires a Log Analytics workspace.

URL:

<https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-nsg-flowlogging-portal>

## Question: 17

You have the Azure virtual networks shown in the following table.

Name	Address space	Subnet	Resource group Azure region
VNet1	10.11.0.0/16	10.11.0.0/17	West US
VNet2	10.11.0.0/17	10.11.0.0/25	West US
VNet3	10.10.0.0/22	10.10.1.0/24	East US
VNet4	192.168.16.0/22	192.168.16.0/24	North Europe

To which virtual networks can you establish a peering connection from VNet1?

- A. VNet2 and VNet3 only
- B. VNet2 only
- C. VNet3 and VNet4 only
- D. VNet2, VNet3, and VNet4

**Answer:**

A

**Explanation:**

Azure virtual network peering requires that the connected virtual networks have non-overlapping IP address spaces. VNet1 has an address space of 10.1.0.0/16.

VNet2 (10.2.0.0/16): The address space does not overlap with VNet1. Peering is possible.

VNet3 (10.3.0.0/16): The address space does not overlap with VNet1. Global VNet peering is possible as they are in different regions.

VNet4 (10.1.0.0/16): The address space is identical to VNet1, causing an overlap. Peering is not possible.

Therefore, VNet1 can only be peered with VNet2 and VNet3.

**Why Incorrect Options are Wrong:**

B. VNet2 only: This is incorrect because Global VNet Peering allows connection to VNet3 in a different region, as its address space is unique.

C. VNet3 and VNet4 only: This is incorrect because VNet4 has an overlapping IP address space with VNet1, which prevents peering.

D. VNet2, VNet3, and VNet4: This is incorrect because VNet4's overlapping IP address space with VNet1 makes peering impossible.

**References:**

Microsoft Learn: Virtual network peering. This official documentation lists the constraints for VNet peering. Under the "Constraints" section, it explicitly states: "The virtual networks must

<https://certempire.com>



have non-overlapping IP address spaces."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-peeringoverview#constraints>

Microsoft Learn: Create, change, or delete a virtual network peering. This tutorial reinforces the requirement. In the prerequisites section, it notes the need for "two virtual networks with non-overlapping IP address spaces."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-managepeering?tabs=peering-portal#prerequisites>

CertEmpire

## Question: 18

You have an Azure subscription that contains a virtual network named VNet1. VNet1 contains four subnets named Gateway, Perimeter, NVA, and Production. The NVA subnet contains two network virtual appliances (NVAs) that will perform network traffic inspection between the Perimeter subnet and the Production subnet. You need to implement an Azure load balancer for the NVAs. The solution must meet the following requirements: The NVAs must run in an active-active configuration that uses automatic failover. The load balancer must load balance traffic to two services on the Production subnet. The services have different IP addresses. Which three actions should you perform? Each correct answer presents part of the solution.

- A. Deploy a basic load balancer
- B. Deploy a standard load balancer
- C. Add two load balancing rules that have HA Ports and Floating IP enabled
- D. Add two load balancing rules that have HA Ports enabled and Floating IP disabled
- E. Add a frontend IP configuration, a backend pool, and a health probe
- F. Add a frontend IP configuration, two backend pools, and a health probe

### Answer:

CertEmpire

B, C, E

### Explanation:

To achieve a highly available active-active configuration for Network Virtual Appliances (NVAs), a Standard Load Balancer is required. The Basic SKU does not support the necessary features. This makes option (B) Deploy a standard load balancer a mandatory first step.

All load balancer configurations require three fundamental components: a frontend IP configuration to receive traffic, a backend pool containing the resources to which traffic is distributed (in this case, the two NVAs), and a health probe to monitor the availability of the backend instances. This makes option (E) Add a frontend IP configuration, a backend pool, and a health probe a correct and necessary action.

For active-active NVA failover, you must configure an HA (High Availability) Ports rule. This single rule type load balances all TCP and UDP flows on all ports. Additionally, Floating IP (also known as Direct Server Return) must be enabled. This ensures traffic symmetry by allowing the NVA to respond directly to the client, which is critical for stateful packet inspection. The requirement to support two distinct services can be met by configuring two frontend IPs on the load balancer, each with its own HA Ports rule, thus requiring (C) two load balancing rules that have HA Ports and Floating IP enabled.

**Why Incorrect Options are Wrong:**

A. Deploy a basic load balancer: The Basic Load Balancer SKU does not support the HA Ports feature, which is essential for the active-active NVA requirement.

D. Add two load balancing rules that have HA Ports enabled and Floating IP disabled: Disabling Floating IP would break the symmetric traffic flow (where request and response traverse the same NVA), which is critical for stateful NVA operation.

F. Add a frontend IP configuration, two backend pools, and a health probe: The two NVAs should be placed in a single backend pool to be load-balanced as a high-availability set. Using two backend pools is incorrect for this scenario.

## References:

1. Standard Load Balancer and HA Ports: Microsoft Azure Documentation. "Load Balancer SKUs". This document explicitly states that HA Ports are only available with the Standard Load Balancer SKU.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/skus>

2. NVA High Availability Configuration: Microsoft Azure Documentation. "Tutorial: Create an NVA and test high availability". This tutorial details the architecture for highly available NVAs, specifying the use of an internal Standard Load Balancer, an HA Ports rule, and enabling Floating IP.

CertEmpire

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/tutorial-nva-load-balancer>

3. Load Balancer Components: Microsoft Azure Documentation. "Azure Load Balancer components". This page describes the essential components of a load balancer, including frontend IP configurations, backend pools, and health probes.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/components>

4. Multiple Frontends for Load Balancer: Microsoft Azure Documentation. "Multiple frontends for Azure Load Balancer". This document explains that a load balancer can have multiple frontend IPs, and "a load-balancing rule is required for each frontend IP." This supports the interpretation that two rules are needed to handle traffic for two distinct services.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/load-balancer-multivip-overview>

## Question: 19

You have an Azure subscription named Subscription1 that contains two Azure virtual networks named VNet1 and VNet2. VNet1 contains a VPN gateway named VPNGW1 that uses static routing. There is a site-to-site VPN connection between your on-premises network and VNet1. On a computer named Client1 that runs Windows 10, you configure a point-to-site VPN connection to VNet1. You configure virtual network peering between VNet1 and VNet2. You verify that you can connect to VNet2 from the on-premises network. Client1 is unable to connect to VNet2. You need to ensure that you can connect Client1 to VNet2. What should you do?

- A. Download and re-install the VPN client configuration package on Client1.
- B. Select Allow gateway transit on VNet1.
- C. Select Allow gateway transit on VNet2.
- D. Enable BGP on VPNGW1

### Answer:

A

### Explanation:

When virtual network peering is configured with **Allow gateway transit**, the routes to the peered VNet (VNet2) are automatically propagated to the VPN gateway (VPNGW1) in VNet1. However, these new routes are not automatically pushed to existing Point-to-Site (P2S) VPN clients. The VPN client configuration package contains the routing information that is installed on the client machine. To update the client's routing table with the address space of VNet2, the administrator must generate a new VPN client configuration package and then download and re-install it on Client1. This action ensures the client has the necessary route to communicate with resources in the peered VNet.

### Why Incorrect Options are Wrong:

- B. Select Allow gateway transit on VNet1: This must already be configured. The scenario states that the on-premises network can connect to VNet2, which requires gateway transit to be enabled on VNet1's peering link.
- C. Select Allow gateway transit on VNet2: The correct setting on the VNet without the gateway (VNet2) is "Use remote gateways," not "Allow gateway transit." This is also implied to be working already.
- D. Enable BGP on VPNGW1: Changing the gateway's routing protocol from static to BGP is unnecessary. The core issue is the client's outdated routing configuration, not the gateway's route propagation mechanism.

**References:**

1. Microsoft Azure Documentation - Configure VPN gateway transit for virtual network peering: This official document explicitly states the requirement for P2S clients. Under the "Point-to-site clients" section, it says, "For point-to-site clients, you must download and reinstall the P2S client package for the routes to the peered virtual network to be advertised to the clients."

URL:

<https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-peering-gatewaytransit#p2s>

2. Microsoft Azure Documentation - About Point-to-Site VPN routing: This document explains how routes are advertised to VPN clients. It clarifies that clients use the routes specified in the VPN client configuration files, which must be updated when the network topology changes.

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/point-to-site-about#routing>

CertEmpire

## Question: 20

### HOTSPOT -

You have an Azure subscription. The subscription contains virtual machines that run Windows Server 2016 and are configured as shown in the following table.

Name	Virtual network	DNS suffix configured in Windows Server
VM1	VNET2	Contoso.com
VM2	VNET2	None
VM3	VNET2	Adatum.com

You create a public Azure DNS zone named adatum.com and a private Azure DNS zone named contoso.com.

You create a virtual network link for contoso.com as shown in the following exhibit.

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/8372f433-2dcd-4361-b5ef-5b188fed87d0/resourceGroups/RG2/provi...

Virtual network

VNET2

Configuration

☒ Enable auto registration ⓘ

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

## 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"/>

Hot Area:

### Answer:

Statement 1: When VM1 starts, a record for VM1 is added to the contoso.com DNS zone.

Yes

Statement 2: When VM2 starts, a record for VM2 is added to the contoso.com DNS zone.

Yes

Statement 3: When VM3 starts, a record for VM3 is added to the adatum.com DNS zone.

CertEmpire

No

### Explanation:

The virtual network VNET2 is linked to the Azure private DNS zone contoso.com with "Enable auto registration" turned on. This configuration automatically creates DNS A records in the contoso.com zone for any virtual machine within VNET2.

1. VM1 and VM2: Both are in VNET2. Therefore, upon starting, Azure automatically registers DNS records for both VM1 and VM2 in the contoso.com zone. The DNS suffix configured within the VM's operating system does not prevent this Azure-level registration feature.

2. VM3: The auto-registration link is exclusively for the contoso.com private zone. The adatum.com zone is a separate, public DNS zone and is not linked to VNET2 for auto-registration. Therefore, no record for VM3 is automatically created in adatum.com.

### References:

1. Microsoft Azure Official Documentation - What is an Azure private DNS zone?: This document details the auto-registration feature. It states, "With the autoregistration feature, whenever you create a new VM in a virtual network that is linked to a private zone, the DNS records for the VM are automatically



created." This supports the "Yes" answers for VM1 and VM2.

o URL: <https://learn.microsoft.com/en-us/azure/dns/private-dns/overview#automatic-registration>

2. Microsoft Azure Official Documentation - Scenarios: Under "Name resolution for VMs within the same virtual network," the documentation explains that by linking a virtual network to a private zone as a registration network, "DNS records for the virtual machines in that virtual network are automatically added to the private zone." This confirms that VMs in VNET2 will register to contoso.com.

o URL: <https://learn.microsoft.com/en-us/azure/dns/private-dns/scenarios#name-resolution-for-vms-within-the-same-virtual-network>

CertEmpire

## Question: 21

You have an Azure subscription that contains the resources in the following table. To which subnets can you apply NSG1?

- A. the subnets on VNet1 only
- B. the subnets on VNet2 and VNet3 only
- C. the subnets on VNet2 only
- D. the subnets on VNet3 only
- E. the subnets on VNet1, VNet2, and VNet3

### Answer:

D

### Explanation:

A Network Security Group (NSG) can only be associated with virtual networks (and their subnets or network interfaces) that exist in the same Azure region. According to the provided table, NSG1 is in the 'East US' region. Of the virtual networks listed, only VNet3 is also in the 'East US' region. Therefore, NSG1 can only be applied to the subnets within VNet3. The resource group of the NSG and the VNet do not need to be the same for the association to be valid.

CertEmpire

### Why Incorrect Options are Wrong:

•

A, C, E: These options are incorrect because VNet1 and VNet2 are in the 'West US' region, while NSG1 is in the 'East US' region. An NSG cannot be associated with resources in a different region.

•

B: This option is incorrect because VNet2 is in a different region ('West US') than NSG1 ('East US').

### References:

•

Microsoft Azure Documentation - Network security groups: "A network security group can be associated to one or more subnets in a virtual network. ... The network security group and the virtual network must exist in the same region."

o Direct URL:

<https://learn.microsoft.com/en-us/azure/virtualnetwork/network-security-group-how-it-works> (This

<https://certempire.com>

principle is detailed

in the "Association of a network security group" section).

- 

Microsoft Azure Documentation - Create, change, or delete a network security group: In the Azure portal or CLI/PowerShell instructions for creating and managing NSGs, the 'Location' (Region) is a mandatory parameter. The documentation on associating this NSG with a subnet implicitly requires both resources to be accessible within the same regional scope.

- o Direct URL:

<https://learn.microsoft.com/en-us/azure/virtualnetwork/manage-network-security-group> (See steps for associating an

NSG with a subnet).

CertEmpire

## Question: 22

### HOTSPOT -

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

Name	Location
RG1	West US
RG2	East US

RG1 contains the resources shown in the following table.

Name	Type	Location
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 running and connects to NIC1 and Disk1. NIC1 connects to VNET1.

RG2 contains a public IP address named IP2 that is in the East US location. IP2 is not assigned to a virtual machine.

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

### Answer Area

Statements	Yes	No
You can move storage1 to RG2.	<input type="radio"/>	<input type="radio"/>
You can move NIC1 to RG2.	<input type="radio"/>	<input type="radio"/>
If you move IP2 to RG1, the location of IP2 will change.	<input type="radio"/>	<input type="radio"/>

Hot Area:

### Answer:

1. You can associate IP2 with NIC1. Yes
2. You can move IP2 to RG1. Yes

3. You can move VNET1 to RG2 without moving any other resource. No

**Explanation:**

A public IP address can be associated with any network interface in the same subscription and region, regardless of the resource-group boundaries. IP2 (East US) therefore can be attached to NIC1 (East US).

Public IP addresses are a supported resource type for the Azure Resource Manager move operation, so IP2 can be moved from RG2 to RG1 as long as both resource groups are in the same subscription and region.

A virtual network cannot be moved by itself when dependent resources (for example, NIC1/VM1) are attached to it; you must move the VNet and all dependent resources together. Therefore, moving only VNET1 to RG2 is not permitted.

**References:**

1. Azure Virtual Network Associate a public IP address to a network interface

<https://learn.microsoft.com/azure/virtual-network/virtual-networks-ip-addresses#associate-a-public-ip-address>

2. Azure Resource Manager Move resources to a new resource group or subscription

<https://learn.microsoft.com/azure/azure-resource-manager/management/move-resourcegroup-and-subscription#supported-resources>

See Public IP address Supported

CertEmpire

See Virtual network Move requires moving dependent resources

## Question: 23

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

- A. Deploy an internal load balancer
- B. Peer VNET1 to another virtual network
- C. Connect webapp1 to VNET1
- D. Deploy an Azure Application Gateway

### Answer:

C

### Explanation:

The core requirement is to enable a connection from an Azure App Service (webapp1) to a resource (VM1) inside a private virtual network (VNET1). The designated Azure feature for this is VNet Integration. By connecting webapp1 to VNET1 using VNet Integration, the web app can make outbound calls to private IP addresses within the VNET. This allows webapp1 to securely access the MySQL database running on VM1's private IP address, resolving the connectivity issue directly and securely without exposing the database to the public internet.

### Why Incorrect Options are Wrong:

- A. Deploy an internal load balancer: An internal load balancer distributes traffic within a VNET. It does not solve the fundamental problem of connecting the external web app to the VNET.
- B. Peer VNET1 to another virtual network: VNet peering connects two distinct virtual networks. Since the web app is a PaaS service and not natively within a VNet, this option is irrelevant.
- D. Deploy an Azure Application Gateway: An Application Gateway is a web traffic manager for inbound requests to web applications. It is not used to facilitate outbound connectivity from a web app to a database.

### References:

1. Correct Answer & Explanation: Microsoft Learn. (2023). Integrate your app with an Azure virtual network. "VNet Integration gives your app access to resources in your virtual network... VNet Integration is used only to make outbound calls from your app into your

virtual network."



URL: <https://learn.microsoft.com/en-us/azure/app-service/overview-vnet-integration>

2. Incorrect Option A: Microsoft Learn. (2024). What is Azure Load Balancer?. "An internal (or private) load balancer is used to distribute traffic to virtual machines and cloud services inside a virtual network."

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/load-balanceroverview#internal-load-balancer>

3. Incorrect Option B: Microsoft Learn. (2024). Azure Virtual Network peering. "Virtual network peering enables you to seamlessly connect two or more Virtual Networks in Azure."

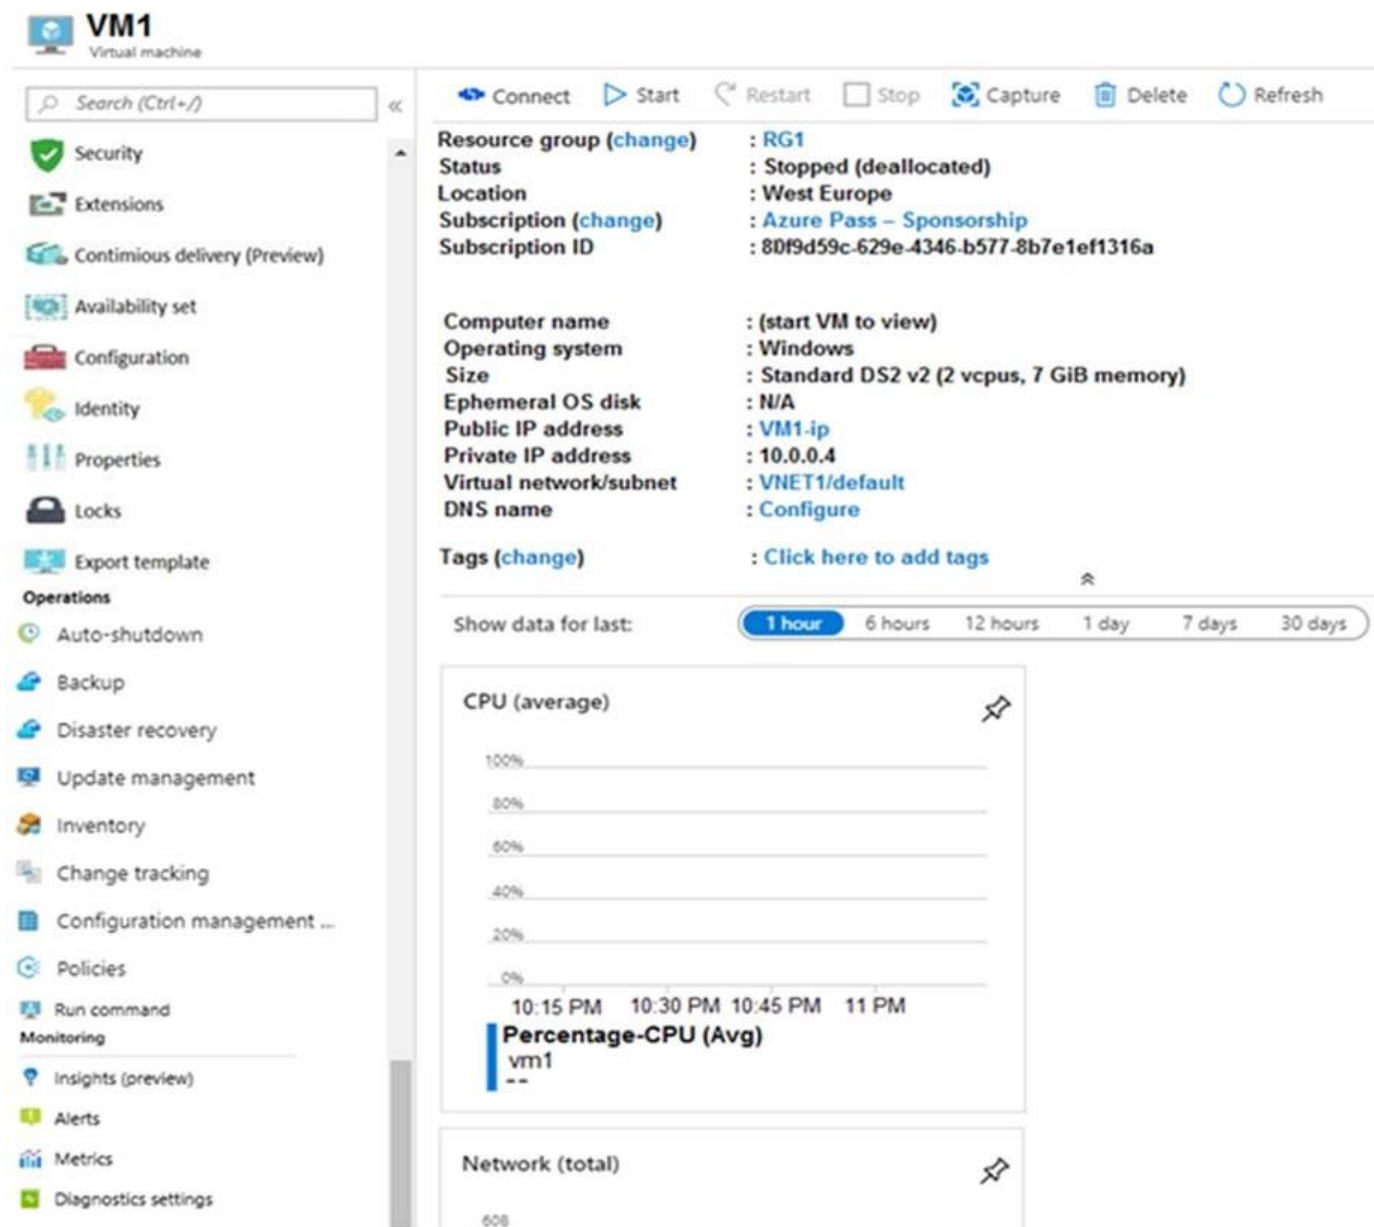
URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-peeringoverview>

4. Incorrect Option D: Microsoft Learn. (2024). What is Azure Application Gateway?. "Azure Application Gateway is a web traffic load balancer that enables you to manage traffic to your web applications."

URL: <https://learn.microsoft.com/en-us/azure/application-gateway/overview>

## Question: 24

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



You need to enable Desired State Configuration for VM1. What should you do first?

- A. Connect to VM1.
- B. Start VM1.
- C. Capture a snapshot of VM1.
- D. Configure a DNS name for VM1.

**Answer:**

B

**Explanation:**

The exhibit shows that the status of VM1 is "Stopped (deallocated)". To enable Azure Automation Desired State Configuration (DSC) on a virtual machine, the Azure VM agent must be running and able to communicate with the Azure control plane. The VM agent only runs when the virtual machine is in a "Running" state. Therefore, the first and most fundamental step required to apply any configuration or extension, including DSC, is to start the virtual machine.

**Why Incorrect Options are Wrong:**

- A. Connect to VM1: You cannot establish a connection (e.g., via RDP) to a virtual machine that is in a "Stopped (deallocated)" state.
- C. Capture a snapshot of VM1: While taking a snapshot is a good practice before making configuration changes, it is not a mandatory prerequisite for enabling DSC.
- D. Configure a DNS name for VM1: A DNS name is used for name resolution over a network and is not a requirement for enabling DSC on the VM itself.

**References:**

Microsoft Azure Documentation - Virtual machine states and billing status: This document explains that in the "Stopped (deallocated)" state, the VM is powered off, and the compute resources are released. Management tasks that require a running OS cannot be performed.

URL: <https://learn.microsoft.com/en-us/azure/virtual-machines/states-billing#states-andbilling>

Microsoft Azure Documentation - Enable Azure Automation State Configuration: The prerequisites for enabling DSC on a VM implicitly require the machine to be running so the DSC extension can be installed and the node can be registered.

URL:

<https://learn.microsoft.com/en-us/azure/automation/automation-dsconboarding#enable-azure-vm>

## Question: 25

You have five Azure virtual machines that run Windows Server 2016. The virtual machines are configured as web servers. You have an Azure load balancer named LB1 that provides load balancing services for the virtual machines. You need to ensure that visitors are serviced by the same web server for each request. What should you configure?

- A. Floating IP (direct server return) to Disabled
- B. Session persistence to None
- C. Floating IP (direct server return) to Enabled
- D. Session persistence to Client IP

### Answer:

D

### Explanation:

To ensure that visitors are consistently serviced by the same web server, you must configure session persistence, also known as source IP affinity. In an Azure Load Balancer rule, setting "Session persistence" to "Client IP" directs all requests from the same source IP address to the same backend virtual machine. This maintains the client's session with a specific server, fulfilling the requirement.

CertEmpire

### Why Incorrect Options are Wrong:

- A. Floating IP (direct server return) to Disabled: Floating IP is unrelated to session persistence. It is a feature used for high availability scenarios, such as with SQL Always On Availability Groups.
- B. Session persistence to None: This is the default setting and explicitly disables session persistence, meaning requests from the same client can be sent to any available backend server.
- C. Floating IP (direct server return) to Enabled: Enabling Floating IP does not configure session affinity. It alters the IP address mapping for specific high-availability patterns, not for maintaining client sessions.

### References:

Azure Load Balancer distribution modes: This official Microsoft documentation explains the different session persistence modes. It states, "Source IP affinity (also called session affinity or client IP affinity) is an Azure Load Balancer distribution mode... This mode uses a 2-tuple (Source IP, Destination IP) or 3-tuple (Source IP, Destination IP, Protocol) hash to map traffic to the available servers."

<https://certempire.com>

Microsoft Learn. (2024). Azure Load Balancer distribution modes. Retrieved from <https://learn.microsoft.com/en-us/azure/load-balancer/load-balancer-distribution-mode>  
Azure Load Balancer Floating IP: This document clarifies the purpose of Floating IP, distinguishing it from session persistence. It describes its use in scenarios like Network Virtual Appliances (NVAs) and SQL Always On.

Microsoft Learn. (2024). Azure Load Balancer Floating IP configuration. Retrieved from <https://learn.microsoft.com/en-us/azure/load-balancer/load-balancer-floating-ip>

CertEmpire

## Question: 26

You have an Azure subscription that contains the following resources:

A virtual network that has a subnet named Subnet1

Two network security groups (NSGs) named NSG-VM1 and NSG-Subnet1

A virtual machine named VM1 that has the required Windows Server configurations to allow Remote Desktop connections

NSG-Subnet1 has the default inbound security rules only.

NSG-VM1 has the default inbound security rules and the following custom inbound security rule:

Priority: 100

Source: Any

Source port range: \*

Destination: \*

Destination port range: 3389

Protocol: UDP

Action: Allow

VM1 has a public IP address and is connected to Subnet1. NSG-VM1 is associated to the network interface of VM1. NSG-Subnet1 is associated to Subnet1.

You need to be able to establish Remote Desktop connections from the internet to VM1.

Solution: You add an inbound security rule to NSG-Subnet1 that allows connections from the Any source to the \*destination for port range 3389 and uses the TCP protocol. You remove NSG-VM1 from the network interface of VM1.

Does this meet the goal?

A. Yes

B. No

### Answer:

A (Yes)

### Explanation:

The proposed solution correctly configures the environment to allow Remote Desktop Protocol (RDP) connections. The initial problem was that the custom rule in NSG-VM1 allowed traffic on port 3389 but specified the UDP protocol, whereas RDP requires the TCP protocol.

The solution resolves this by adding a new inbound security rule to NSG-Subnet1 that correctly allows TCP traffic on port 3389. By then removing the incorrectly configured NSG-VM1 from the virtual machine's network interface, the only active NSG (NSG-Subnet1) now contains the necessary rule to permit RDP traffic from the internet, thus meeting the goal.

**References:**

1. Network Security Groups: Microsoft Learn. "Network security groups filter network traffic to and from Azure resources in an Azure virtual network. An NSG contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure resources... For inbound traffic, Azure processes the rules in a network security group associated to a subnet first, if there is one, and then the rules in a network security group associated to the network interface."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groupsoverview>

2. Create, change, or delete a network security group: Microsoft Learn. "To allow RDP to a Windows VM, you can add a security rule to a network security group. You can add the rule to an existing NSG or when you create a VM... By default, RDP is allowed on TCP port 3389."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/manage-network-securitygroup?tabs=network-security-group-portal#create-a-security-rule>

3. Troubleshoot Remote Desktop connections to an Azure virtual machine: Microsoft Learn. "Check network security group rules... Check that an inbound rule for the TCP port 3389 exists in the NSG that's associated with both the network interface of the VM and the subnet that it's in."

CertEmpire

URL:  
<https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshootrdp-connection>



## Question: 27

You have an Azure subscription that contains the following resources: A virtual network that has a subnet named Subnet1 Two network security groups (NSGs) named NSG-VM1 and NSG-Subnet1 A virtual machine named VM1 that has the required Windows Server configurations to allow Remote Desktop connections NSG-Subnet1 has the default inbound security rules only. NSG-VM1 has the default inbound security rules and the following custom inbound security rule: Priority: 100 Source: Any Source port range: \* Destination: \* Destination port range: 3389 Protocol: UDP - Action: Allow VM1 has a public IP address and is connected to Subnet1. NSG-VM1 is associated to the network interface of VM1. NSG-Subnet1 is associated to Subnet1. You need to be able to establish Remote Desktop connections from the internet to VM1. Solution: You add an inbound security rule to NSG-Subnet1 that allows connections from the internet source to the VirtualNetwork destination for port range 3389 and uses the UDP protocol. Does this meet the goal?

A. Yes

B. No

### Answer:

B

CertEmpire

### Explanation:

The solution does not meet the goal because it configures a rule for the incorrect protocol. The Remote Desktop Protocol (RDP) requires the TCP protocol on port 3389 for communication. The proposed solution adds an inbound security rule to NSG-Subnet1 that allows traffic on port 3389 using the UDP protocol.

Since the RDP client will initiate a connection using TCP, this traffic will not match the new UDP rule. Consequently, the traffic will be evaluated against the default DenyAllInbound rule in NSG-Subnet1 and be blocked. The connection will fail before it even reaches the network interface NSG (NSG-VM1).

### References:

1. Microsoft Azure Documentation - Network Security Groups: "Network security groups use security rules to filter network traffic... For each rule, you can specify a source and destination, port, and protocol." This documentation confirms that the protocol is a critical and distinct component of a security rule.

URL: <https://docs.microsoft.com/en-us/azure/virtual-network/network-security-groupsoverview#security-rules>

2. Microsoft Azure Documentation - Troubleshoot Remote Desktop connections to an Azure virtual machine: This official troubleshooting guide explicitly states the port and protocol

required for RDP. "By default, RDP connections are on TCP port 3389."

URL: <https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshootrdp-connection#check-nsg-rules>

3. Microsoft Azure Documentation - How network security groups filter network traffic: "For inbound traffic, Azure processes the rules in a network security group associated to a subnet first, if there is one, and then the rules in a network security group associated to the network interface." This confirms the order of processing, showing the connection would be blocked at the subnet level.

URL: <https://docs.microsoft.com/en-us/azure/virtual-network/network-security-group-how-itworks#inbound-traffic>

CertEmpire

## Question: 28

You have an Azure subscription that contains the following resources: A virtual network that has a subnet named Subnet1 Two network security groups (NSGs) named NSG-VM1 and NSG-Subnet1 A virtual machine named VM1 that has the required Windows Server configurations to allow Remote Desktop connections NSG-Subnet1 has the default inbound security rules only. NSG-VM1 has the default inbound security rules and the following custom inbound security rule: Priority: 100 Source: Any Source port range: \* Destination: \* Destination port range: 3389 Protocol: UDP Action: Allow VM1 has a public IP address and is connected to Subnet1. NSG-VM1 is associated to the network interface of VM1. NSG-Subnet1 is associated to Subnet1. You need to be able to establish Remote Desktop connections from the internet to VM1. Solution: You add an inbound security rule to NSG-Subnet1 and NSG-VM1 that allows connections from the internet source to the VirtualNetwork destination for port range 3389 and uses the TCP protocol. Does this meet the goal?

A. Yes

B. No

### Answer:

A. Yes

CertEmpire

### Explanation:

The goal is to enable Remote Desktop Protocol (RDP) connections, which require inbound traffic on TCP port 3389. In this scenario, network traffic to VM1 is filtered by two Network Security Groups (NSGs): first by NSG-Subnet1 at the subnet level, and then by NSG-VM1 at the network interface (NIC) level. For the connection to succeed, both NSGs must allow the traffic.

The proposed solution adds a new inbound rule to both NSGs to explicitly allow TCP traffic on port 3389 from the internet. This correctly addresses the filtering at both the subnet and NIC layers, ensuring the RDP connection is permitted and meeting the goal.

### References:

1. Microsoft Documentation: How network security groups filter network traffic. This document explains that for inbound traffic, rules in the subnet's NSG are processed first, followed by rules in the NIC's NSG. Both must allow the traffic.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-group-how-itworks#inbound-traffic>

2. Microsoft Documentation: Network security group security rules. This page details the properties of a security rule, including protocol (TCP, UDP, etc.), port ranges, and action

(Allow/Deny). It confirms that separate rules are needed for different protocols.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groupsoverview#security-rules>

3. Microsoft Documentation: Troubleshoot Remote Desktop connections to an Azure virtual machine. This guide explicitly states that for RDP to work, an inbound security rule allowing TCP traffic on port 3389 is required.

URL: <https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshootrdp-connection#check-nsg-rules>

CertEmpire

## Question: 29

### HOTSPOT -

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

```

Name                : VNet1
ResourceGroupName   : Production
Location            : westus
Id                  : /subscriptions/14d26092-8e42-4ea7-b770-
9dcef70fb1ea/resourceGroups/Production/providers/Microsoft.Network/virtualNetworks/VNet1
Etag                : W/"76f7edd6-d022-455b-aeae-376059318e5d"
ResourceGuid        : 562696cc-b2ba-4cc5-9619-0a735d6c34c7
ProvisioningState    : Succeeded
Tags                :
AddressSpace        : {
                      "AddressPrefixes": [
                        "10.2.0.0/16"
                      ]
                      }
DhcpOptions          : {}
Subnets             : [
                      {
                        "Name": "default",
                        "Etag": "W/\ "76f7edd6-d022-455b-aeae-376059318e5d\\"",
                        "Id": "/subscriptions/14d26092-8e42-4ea7-b770-
9dcef70fb1ea/resourceGroups/Production/providers/Microsoft.Network/
virtualNetworks/VNet1/subnets/default",
                        "AddressPrefix": "10.2.0.0/24",
                        "IpConfigurations": [],
                        "ResourceNavigationLinks": [],
                        "ServiceEndpoints": [],
                        "ProvisioningState": "Succeeded"
                      }
                      ]
VirtualNetworkPeerings : []
EnableDDoSProtection  : false
EnableVmProtection     : false

```

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

Hot Area:

## Answer Area

Before a virtual machine on VNet1 can receive an IP address from 192.168.1.0/24, you must first

	▼
add a network interface	
add a subnet	
add an address space	
delete a subnet	
delete an address space	

Before a virtual machine on VNet1 can receive an IP address from 10.2.1.0/24, you must first

	▼
add a network interface	
add a subnet	
add an address space	
delete a subnet	
delete an address space	

### Answer:

Before a virtual machine on VNet1 can receive an IP address from 192.168.1.0/24, you must first add an address space. CertEmpire

Before a virtual machine on VNet1 can receive an IP address from 10.2.1.0/24, you must first add a subnet.

### Explanation:

1. To use the 192.168.1.0/24 range, you must first add an address space. The VNet's current and only address space is 10.2.0.0/16. Any IP address or subnet created for this VNet must be within this space. Since 192.168.1.0/24 is outside of the 10.2.0.0/16 range, you must first add a new address space (e.g., 192.168.0.0/16) to the VNet's configuration. Only after adding the new address space can you create a subnet within it.
2. To use the 10.2.1.0/24 range, you must first add a subnet. This IP range is already contained within the VNet's existing 10.2.0.0/16 address space. However, virtual machines receive their IP addresses from a subnet, not directly from the VNet's address space. The only existing subnet is 10.2.0.0/24. Therefore, a new subnet with the 10.2.1.0/24 prefix must be created before a VM can be assigned an IP from that range.

### References:

1. Microsoft Learn: Azure Virtual Network concepts and best practices.
  - o Quote/Concept: "When you create a virtual network, you must specify



one or more custom private IP address ranges using CIDR notation...

The address ranges you specify can't overlap with any other address ranges of the virtual networks that you've connected in your subscription and your on-premises networks... You can add subnets to your virtual network after you create it." This supports the principle that all subnets must fall within a defined VNet address space.

o URL: <https://learn.microsoft.com/en-us/azure/virtual-network/conceptsand-best-practices> (Section: Address space)

2. Microsoft Learn: Add, change, or remove a virtual network address space.

o Quote/Concept: "If you need to add address spaces... You can add IP address spaces to a virtual network at any time... The address spaces you add can't overlap with the other address spaces for the virtual network or with your on-premises networks." This directly supports the answer for the 192.168.1.0/24 scenario.

o URL: <https://learn.microsoft.com/en-us/azure/virtual-network/managevirtual-network#add-or-remove-an-address-space>

3. Microsoft Learn: Add, change, or delete a virtual network subnet. CertEmpire

o Quote/Concept: "A subnet is a range of IP addresses in your virtual network... The address range for the subnet must be unique within the address space of the virtual network and can't overlap with other subnet address ranges in the virtual network." This supports the answer for the 10.2.1.0/24 scenario, as a new subnet must be created within the existing address space.

o URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtualnetwork-manage-subnet#add-a-subnet>

## Question: 30

You have an Azure subscription that contains a virtual network named VNET1. VNET1 contains the subnets shown in the following table.

Name	Connected virtual machines
Subnet1	VM1, VM2
Subnet2	VM3, VM4
Subnet3	VM5, VM6

Each virtual machine uses a static IP address.

You need to create network security groups (NSGs) to meet following requirements:

Allow web requests from the internet to VM3, VM4, VM5, and VM6.

Allow all connections between VM1 and VM2.

Allow Remote Desktop connections to VM1.

Prevent all other network traffic to VNET1.

What is the minimum number of NSGs you should create?

- A. 1
- B. 3
- C. 4
- D. 12

CertEmpire

**Answer:**

B

**Explanation:**

To meet the specified requirements with the minimum number of Network Security Groups (NSGs), we must group resources with identical security policies.

1. NSG for Web Servers: The requirement to allow web requests is identical for VM3, VM4, VM5, and VM6. These VMs are in Subnet2 and Subnet3. A single NSG can be created with an inbound rule to allow HTTP/HTTPS traffic and be associated with both Subnet2 and Subnet3. This accounts for the first NSG.

2. NSGs for Subnet1: Within Subnet1, VM1 and VM2 have different requirements. VM1 requires Remote Desktop access, while VM2 does not. To apply different rules to VMs within the same subnet, we must use NSGs at the Network Interface (NIC) level. Therefore, VM1 needs its own NSG to allow RDP and traffic from VM2. VM2 needs a separate NSG to allow traffic from VM1.

This results in a minimum of three NSGs: one for the web subnets, one for VM1's NIC, and one for VM2's NIC.

### Why Incorrect Options are Wrong:

- A. 1: A single NSG cannot satisfy the conflicting requirements (e.g., allowing RDP to VM1 but not VM2, and allowing web traffic to other specific VMs).
- C. 4: This is not the minimum. The web servers in Subnet2 and Subnet3 have identical requirements and can share a single NSG. Creating four NSGs is unnecessary.
- D. 12: This is excessive. Creating an NSG for every NIC and every subnet is inefficient and not required to meet the security policies.

### References:

Microsoft Documentation: Network security groups. This document explains how NSGs work and how they are associated. It states, "You can associate zero, or one, network security group to each virtual network subnet and network interface... If you need to apply different security rules to some virtual machines within a subnet, you can associate a network security group to the network interfaces of those virtual machines." This directly supports the need for separate NSGs for VM1 and VM2.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groupsoverview>

Microsoft Documentation: How network security groups filter network traffic. This page details the rule evaluation process, including the interaction between subnet and NIC NSGs, and the use of default rules. The need to create specific rules to override default behaviors (like AllowVnetInBound) and implement a "deny all other" policy is central to this scenario.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-group-how-itworks>

## Question: 31

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

Name	Type	Resource group
VNET1	Virtual network	RG1
VM1	Virtual machine	RG1

The Not allowed resource types Azure policy that has policy enforcement enabled is assigned to RG1 and uses the following parameters:

Microsoft.Network/virtualNetworks

Microsoft.Compute/virtualMachines

In RG1, you need to create a new virtual machine named VM2, and then connect VM2 to VNET1.

What should you do first?

- A. Remove Microsoft.Compute/virtualMachines from the policy.
- B. Create an Azure Resource Manager template
- C. Add a subnet to VNET1.
- D. Remove Microsoft.Network/virtualNetworks from the policy.

CertEmpire

**Answer:**

A

**Explanation:**

The Azure Policy assigned to the resource group RG1 has the "Not allowed resource types" definition with Microsoft.Compute/virtualMachines listed as a parameter. With policy enforcement enabled, this configuration explicitly blocks the creation of any new virtual machine resources within RG1. To create VM2, which is a Microsoft.Compute/virtualMachines resource, you must first modify the policy assignment to remove this restriction. This makes removing the specific resource type from the policy the necessary initial step before any deployment can succeed.

**Why Incorrect Options are Wrong:**

- B. Create an Azure Resource Manager template: The deployment method (ARM template, Portal, CLI) is irrelevant; the Azure Policy will block the resource creation regardless of how it is initiated.
- C. Add a subnet to VNET1: While a VM requires a subnet, the policy will prevent the VM creation from even starting. The policy block must be resolved first.
- D. Remove Microsoft.Network/virtualNetworks from the policy: This would allow the creation

<https://certempire.com>

of new virtual networks, but it does not address the policy rule that is blocking the creation of the virtual machine.

## References:

1. Azure Policy Effects: The "Deny" effect, used by the "Not allowed resource types" policy, evaluates a resource creation or update request. If the condition is met, the request fails. This is the core mechanism blocking the VM creation.

Source: Microsoft Azure Documentation, "Understand Azure Policy effects - Deny".

URL: <https://learn.microsoft.com/en-us/azure/governance/policy/concepts/effects#deny>

2. Remediate Non-compliant Resources: To create a resource that is currently blocked by a "Deny" policy, you must either modify the policy definition/assignment or create a policy exemption for that specific resource. The most direct action described in the options is to modify the policy assignment.

Source: Microsoft Azure Documentation, "Remediate non-compliant resources with Azure Policy".

URL: <https://learn.microsoft.com/en-us/azure/governance/policy/how-to/remediate-resources>

3. "Not allowed resource types" Policy Definition: This built-in policy definition restricts which resource types can be deployed. The parameters specify the list of types to block.

Source: Microsoft Azure Documentation, "Azure Policy built-in definitions".

URL: <https://learn.microsoft.com/en-us/azure/governance/policy/samples/built-in-policies#general> (Search for "Not allowed resource types" on the page).

## Question: 32

Your company has an Azure subscription named Subscription1. The company also has two on-premises servers named Server1 and Server2 that run Windows Server 2016. Server1 is configured as a DNS server that has a primary DNS zone named adatum.com. Adatum.com contains 1,000 DNS records. You manage Server1 and Subscription1 from Server2.

Server2 has the following tools installed: The DNS Manager console Azure PowerShell Azure CLI 2.0 You need to move the adatum.com zone to an Azure DNS zone in Subscription1. The solution must minimize administrative effort. What should you use?

- A. Azure CLI
- B. Azure PowerShell
- C. the Azure portal
- D. the DNS Manager console

### Answer:

A

### Explanation:

The most efficient method to migrate a large DNS zone with 1,000 records from an on-premises server to Azure is by using the Azure CLI. The Azure CLI provides a specific command, `az network dns zone import`, which is designed to take a standard DNS zone file and automatically create all the corresponding records in an Azure DNS zone. This single-command operation significantly minimizes administrative effort compared to manually creating records or scripting the process.

### Why Incorrect Options are Wrong:

- B. Azure PowerShell: While Azure PowerShell can manage Azure DNS, it lacks a single, dedicated cmdlet for importing an entire zone file. Accomplishing this would require writing a custom script to parse the file and create each record individually, which is more effort.
- C. the Azure portal: Using the Azure portal would require manually creating each of the 1,000 DNS records. This is the most time-consuming and error-prone method, representing the maximum administrative effort.
- D. the DNS Manager console: This tool is used to manage on-premises Windows Server DNS. It has no native capability to interact with or migrate zones directly to the Azure DNS service.

### References:

1. Microsoft Azure Documentation, "Tutorial: Import and export a DNS zone file using the

Azure CLI": This official tutorial explicitly details the use of the az network dns zone import

command as the primary method for importing a zone file. It states, "This article explains how to import and export a DNS zone file for Azure DNS by using the Azure CLI."

URL: <https://learn.microsoft.com/en-us/azure/dns/dns-import-export>

2. Microsoft Azure Documentation, "az network dns zone import": The reference for the specific Azure CLI command confirms its purpose: "Create a DNS zone and records from a local zone file." This directly addresses the requirement to move the zone with minimal effort.

URL: <https://learn.microsoft.com/en-us/cli/azure/network/dns/zone?view=azure-clilatest#az-network-dns-zone-import>

3. Microsoft Azure Documentation, "New-AzDnsRecordSet": The documentation for the comparable Azure PowerShell cmdlet shows it is designed to create individual record sets, confirming that a bulk import requires additional scripting, unlike the Azure CLI's direct import command.

URL: <https://learn.microsoft.com/en-us/powershell/module/az.dns/new-azdnsrecordset>

CertEmpire



## Question: 33

You have a public load balancer that balances ports 80 and 443 across three virtual machines named VM1, VM2, and VM3. You need to direct all the Remote Desktop Protocol (RDP) connections to VM3 only. What should you configure?

- A. an inbound NAT rule
- B. a new public load balancer for VM3
- C. a frontend IP configuration
- D. a load balancing rule

### Answer:

A

### Explanation:

An inbound NAT rule is the correct configuration for forwarding traffic from a specific port on the load balancer's public IP address to a specific virtual machine in the backend pool. The requirement is to direct all Remote Desktop Protocol (RDP) traffic, which uses a specific port (typically 3389), to a single, designated virtual machine (VM3). A load balancing rule, in contrast, distributes traffic across all available VMs in the backend pool, which is not the desired outcome.

CertEmpire

### Why Incorrect Options are Wrong:

- B. a new public load balancer for VM3: This is an inefficient and unnecessary solution. The existing load balancer is capable of handling this requirement with the correct rule configuration, avoiding additional cost and complexity.
- C. a frontend IP configuration: This defines the public IP address for the load balancer. While necessary for the load balancer to function, it does not define the rules for directing traffic to backend resources.
- D. a load balancing rule: A load balancing rule would distribute RDP connections across all three virtual machines (VM1, VM2, and VM3), not direct them exclusively to VM3 as required.

### References:

1. Microsoft Azure Documentation - What is Azure Load Balancer?: "An inbound NAT rule forwards incoming traffic sent to the frontend IP address and port combination to a specific virtual machine or instance in the backend pool. A load balancing rule distributes incoming traffic across all instances within the backend pool."

URL:

<https://certempire.com>

<https://docs.microsoft.com/en-us/azure/load-balancer/load-balanceroverview#inbound-nat-rule>

2. Microsoft Azure Documentation - Manage inbound NAT rules for Azure Load Balancer:  
"Azure Load Balancer supports inbound network address translation (NAT) rules. You use these rules to specify a backend resource to route traffic to from the load balancer frontend."

URL: <https://docs.microsoft.com/en-us/azure/load-balancer/manage-inbound-nat-rules>

CertEmpire

**Question: 34****HOTSPOT -**

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

Name	Subnets
VNet1	Subnet11, Subnet12
VNet2	Subnet13

Subscription1 contains the virtual machines in the following table.

Name	Subnet	Availability set
VM1	Subnet11	AS1
VM2	Subnet11	AS1
VM3	Subnet11	<i>Not applicable</i>
VM4	Subnet11	<i>Not applicable</i>
VM5	Subnet12	<i>Not applicable</i>
VM6	Subnet12	<i>Not applicable</i>

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

Name: LB1

SKU: Basic

Type: Internal

Subnet: Subnet12

Virtual network: VNET1

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

Hot Area:

## Answer Area

Statements	Yes	No
LB1 can balance the traffic between VM1 and VM2.	<input type="radio"/>	<input type="radio"/>
LB1 can balance the traffic between VM3 and VM4.	<input type="radio"/>	<input type="radio"/>
LB1 can balance the traffic between VM5 and VM6.	<input type="radio"/>	<input type="radio"/>

### Answer:

Statements

Yes No

LB1 can balance the traffic between VM1 and VM2

☐

LB1 can balance the traffic between VM3 and VM4.

☐

LB1 can balance the traffic between VM5 and VM6.

☐

CertEmpire

### Explanation:

The key to this question is understanding the backend pool limitations of the Basic SKU Azure Load Balancer. While all the virtual machines are within the same Virtual Network (VNet) as the load balancer (VNET1), the composition of the backend pool is restricted.

1. LB1 can balance the traffic between VM1 and VM2: Yes

o VM1 and VM2 are in the same virtual network (VNET1) as the load balancer.

o Both VM1 and VM2 are members of the same availability set (AS1).

o A Basic Load Balancer's backend pool can be populated by all virtual machines within a single availability set. This is a valid configuration.

2. LB1 can balance the traffic between VM3 and VM4: No

o VM3 and VM4 are standalone virtual machines, as they are not part of any availability set.

o A significant limitation of the Basic Load Balancer is that its backend pool cannot contain more than one standalone virtual machine. Since balancing traffic between VM3 and VM4 would require adding both to the same backend pool, this configuration is not possible.

3. LB1 can balance the traffic between VM5 and VM6: No

- o Similar to the previous statement, VM5 and VM6 are standalone virtual machines.
- o Due to the Basic SKU limitation, a backend pool cannot be configured with two or more standalone virtual machines. Therefore, LB1 cannot balance traffic between VM5 and VM6.

## References:

- 

Microsoft Azure Documentation - Load Balancer SKUs: This document explicitly details the differences between Basic and Standard SKU load balancers. In the feature comparison table, under "Backend pool," it specifies that the Basic SKU backend pool is limited to a "single availability set, single virtual machine scale set, or a single virtual machine." This confirms that multiple standalone VMs are not supported in a single backend pool for the Basic SKU.

o URL: <https://docs.microsoft.com/en-us/azure/load-balancer/skus> (Refer to the "SKU comparison" section).

- 

CertEmpire

Microsoft Azure Documentation - Load Balancer Components: This page details the components of a load balancer, including the backend pool. It states, "For a Basic load balancer, the backend pool can't include more than one virtual machine that isn't part of an availability set or a virtual machine scale set."

o URL: <https://learn.microsoft.com/en-us/azure/loadbalancer/components#backend-pools> (Refer to the "Backend pools" section).

## Question: 35

### HOTSPOT -

You have an Azure virtual machine that runs Windows Server 2019 and has the following configurations:

Name: VM1

Location: West US

Connected to: VNET1

Private IP address: 10.1.0.4

Public IP addresses: 52.186.85.63

DNS suffix in Windows Server: Adatum.com

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

Name	Type	Location
Adatum.pri	Private	West Europe
Contoso.pri	Private	Central US
Adatum.com	Public	West Europe
Contoso.com	Public	North Europe

You need to identify which DNS zones you can link to VNET1 and the DNS zones to which VM1 can automatically register.

CertEmpire

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

Hot Area:

### 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:**

DNS zones that can be linked to VNET1: - adatum.com (Azure Private DNS zone) -



contoso.com (Azure Private DNS zone) - 0.1.10.in-addr.arpa (Azure Private DNS reverse zone) DNS zones to which VM1 will automatically register: - adatum.com (Azure Private DNS zone) - 0.1.10.in-addr.arpa (Azure Private DNS reverse zone)

### Explanation:

Only Azure Private DNS zones can be linked to a virtual network; public zones cannot. Any private zone's forward or reverse can therefore be linked to VNET1.

Automatic (dynamic) registration occurs only for private zones that are linked to the virtual network with registration enabled and whose names match either:

1. the VM DNS suffix (forward lookup), or
2. the appropriate reverse-lookup namespace for the VM's IP address.

VM1 DNS suffix is adatum.com, so it auto-registers in the private adatum.com zone. Its IP 10.1.0.4 falls into the 0.1.10.in-addr.arpa reverse zone, so a PTR record is also auto-registered there. contoso.com does not match the DNS suffix, so no auto-registration occurs in that zone.

### References:

1. Microsoft Azure Docs What is Azure Private DNS? (Functions: linking, auto-registration)

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

2. Microsoft Azure Docs Create a private DNS zone and link it to a virtual network (linking rules, auto-registration option)

<https://learn.microsoft.com/en-us/azure/dns/private-dns-getstarted-portal#link-a-virtual-network-to-the-private-zone>

3. Microsoft Azure Docs Reverse DNS for Private DNS (automatic PTR registration)

<https://learn.microsoft.com/en-us/azure/dns/private-dns-reverse-zones>

## Question: 36

DRAG DROP -

You have an Azure subscription that contains two virtual networks named VNet1 and VNet2. Virtual machines connect to the virtual networks.

The virtual networks have the address spaces and the subnets configured as shown in the following table.

Virtual network	Address space	Subnet	Peering
VNet1	10.1.0.0/16	10.1.0.0/24 10.1.1.0/26	VNet2
VNet2	10.2.0.0/16	10.2.0.0/24	VNet1

You need to add the address space of 10.33.0.0/16 to VNet1. The solution must ensure that the hosts on VNet1 and VNet2 can communicate.

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

Select and Place:

### Actions

Remove VNet1.

Add the 10.33.0.0/16 address space to VNet1.

Create a new virtual network named VNet1.

On the peering connection in VNet2, allow gateway transit.

Recreate peering between VNet1 and VNet2.

On the peering connection in VNet1, allow gateway transit.

Remove peering between VNet1 and VNet2.

### Answer Area



**Answer:**

Actions		Answer Area	
Remove VNet1.		Remove peering between VNet1 and VNet2.	
Add the 10.33.0.0/16 address space to VNet1.		Add the 10.33.0.0/16 address space to VNet1.	
Create a new virtual network named VNet1.	➤	Recreate peering between VNet1 and VNet2.	⬆
On the peering connection in VNet2, allow gateway transit.	⬅		⬇
Recreate peering between VNet1 and VNet2.			
On the peering connection in VNet1, allow gateway transit.			
Remove peering between VNet1 and VNet2.			

### Explanation:

The NSG already contains an outbound rule (DenyWebSites) that denies TCP port 80. Because the NSG is presently attached only to a single NIC, VM2 is unaffected. Associating the NSG with Subnet1 makes every NIC in that subnet including those of VM1 and VM2 subject to the existing outbound-deny rule, preventing both VMs from reaching Internet web sites on port 80.

### References:

1. Microsoft Azure Documentation Network security groups overview, Associations section:  
An NSG linked to a subnet applies to all network interfaces in that subnet.  
<https://learn.microsoft.com/azure/virtual-network/network-security-groupsoverview#associations>
2. Microsoft Azure Documentation Security rules table: Outbound rules filter traffic leaving the VM to the Internet on specified ports.  
<https://learn.microsoft.com/azure/virtualnetwork/network-security-groups-overview#security-rules>

**Question: 37**

DRAG DROP -

You have an on-premises network that you plan to connect to Azure by using a site-to-site VPN.

In Azure, you have an Azure virtual network named VNet1 that uses an address space of 10.0.0.0/16. VNet1 contains a subnet named Subnet1 that uses an address space of 10.0.0.0/24.

You need to create a site-to-site VPN to Azure.

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

NOTE: More than one order of answer choice is correct. You will receive credit for any of the correct orders you select.

Select and Place:

**Actions**

Create a local gateway.

Create a VPN gateway.

Create a gateway subnet.

Create a custom DNS server.

Create a VPN connection.

Create an Azure Content Delivery Network (CDN) profile.

**Answer Area**

**Answer:**

**Actions****Answer Area**

Create a local gateway.

Create a gateway subnet.

Create a VPN gateway.

Create a VPN gateway.

Create a gateway subnet.



Create a local gateway.



Create a custom DNS server.



Create a VPN connection.



Create a VPN connection.

Create an Azure Content Delivery Network (CDN) profile.

**Explanation:**

The NSG already contains an outbound rule (DenyWebSites) that denies TCP port 80.

Because the NSG is presently attached only to a single NIC, VM2 is unaffected.

Associating the NSG with Subnet1 makes every NIC in that subnet including those of VM1 and VM2 subject to the existing outbound-deny rule, preventing both VMs from reaching Internet web sites on port 80.

**References:**

1. Microsoft Azure Documentation Network security groups overview, Associations section:

An NSG linked to a subnet applies to all network interfaces in that subnet.

<https://learn.microsoft.com/azure/virtual-network/network-security-groupsoverview#associations>

2. Microsoft Azure Documentation Security rules table: Outbound rules filter traffic leaving the VM to the Internet on specified ports.

<https://learn.microsoft.com/azure/virtualnetwork/network-security-groups-overview#security-rules>

## Question: 38

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

Name	Type	Details
VNet1	Virtual network	<i>Not applicable</i>
Subnet1	Subnet	Hosted on VNet1
VM1	Virtual machine	On Subnet1
VM2	Virtual machine	On Subnet1

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

Move Delete Refresh

Resource group (change) : RG1lod9053488  
Location : East US  
Subscription (change) : Microsoft AZ  
Subscription ID : ac344a74-f85a-4b2e-8057-642088faaf20  
Tags (change) : Click here to add tags

Custom security rules : 1 inbound, 1 outbound  
Associated with : 0 subnets, 0 network interfaces

PRIORITY	NAME	PORT	PROTOCOL	SOURCE	DESTINATION	ACTION
100	Port_80	80	TCP	Internet	Any	Deny
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	Allow AzureLoadBalancerInBound	Any	Any	AzureLoadBalancer	Any	Allow
65500	DenyAllInBound	Any	Any	Any	Any	Deny

PRIORITY	NAME	PORT	PROTOCOL	SOURCE	DESTINATION	ACTION
100	DenyWebSites	80	TCP	Any	Internet	Deny
65000	AllowVnetOutBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	AllowInternetOutBound	Any	Any	Any	Internet	Allow
65500	DenyAllOutBound	Any	Any	Any	Any	Deny

You need to prevent users of VM1 and VM2 from accessing websites on the Internet over TCP port 80.

What should you do?

- A. Disassociate the NSG from a network interface
- B. Change the Port\_80 inbound security rule.
- C. Associate the NSG to Subnet1.
- D. Change the DenyWebSites outbound security rule.

**Answer:**

C

**Explanation:**

The NSG already contains an outbound rule (DenyWebSites) that denies TCP port 80.

Because the NSG is presently attached only to a single NIC, VM2 is unaffected.

Associating the NSG with Subnet1 makes every NIC in that subnet including those of VM1 and VM2 subject to the existing outbound-deny rule, preventing both VMs from reaching Internet web sites on port 80.

**Why Incorrect Options are Wrong:**

A. Disassociating the NSG removes all its rules, permitting not blocking Internet access.

B. Inbound rules govern traffic entering the VM; they do not control outbound traffic to websites.

D. The outbound deny rule already blocks port 80; no modification is required its scope must simply include both VMs.

**References:**

1. Microsoft Azure Documentation Network security groups overview, Associations section:

An NSG linked to a subnet applies to all network interfaces in that subnet.

<https://learn.microsoft.com/azure/virtual-network/network-security-groupsoverview#associations>

2. Microsoft Azure Documentation Security rules table: Outbound rules filter traffic leaving the VM to the Internet on specified ports.

CertEmpire

<https://learn.microsoft.com/azure/virtualnetwork/network-security-groups-overview#security-rules>



## Question: 39

You have two subscriptions named Subscription1 and Subscription2. Each subscription is associated to a different Azure AD tenant. Subscription1 contains a virtual network named VNet1. VNet1 contains an Azure virtual machine named VM1 and has an IP address space of 10.0.0.0/16. Subscription2 contains a virtual network named VNet2. VNet2 contains an Azure virtual machine named VM2 and has an IP address space of 10.10.0.0/24. You need to connect VNet1 to VNet2. What should you do first?

- A. Move VM1 to Subscription2.
- B. Move VNet1 to Subscription2.
- C. Modify the IP address space of VNet2.
- D. Provision virtual network gateways.

### Answer:

D

### Explanation:

The two virtual networks (VNETs) are in different subscriptions associated with different Azure AD tenants. Standard VNet peering cannot connect VNETs across different tenants. Therefore, a VNet-to-VNet VPN gateway connection is required. The foundational and first step to establish this type of connection is to create a virtual network gateway in each VNet. The existing IP address spaces do not overlap, which is a prerequisite for this connection, so no changes are needed there.

### Why Incorrect Options are Wrong:

- A. Move VM1 to Subscription2: Moving a virtual machine does not establish connectivity between the virtual networks themselves.
- B. Move VNet1 to Subscription2: This is a significant administrative change. The direct method to connect the networks as they are is with gateways, not by moving resources between tenants.
- C. Modify the IP address space of VNet2: The IP address spaces (10.0.0.0/16 and 10.10.0.0/24) do not overlap, so modification is unnecessary.

### References:

1. Microsoft Azure Documentation - Configure a VNet-to-VNet VPN gateway connection by using the Azure portal: This official guide outlines the procedure. The first major configuration step after creating the VNETs is to "Create the virtual network gateways."

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-howto-vnet-vnetresource-manager-portal#create-the-virtual-network-gateways>

<https://certempire.com>



2. Microsoft Azure Documentation - Virtual network peering: This document clarifies the limitations of VNet peering, stating that while it can work across subscriptions, those subscriptions must be associated with the same Azure Active Directory tenant. This confirms peering is not an option in the given scenario.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-peeringoverview#requirements-and-constraints>

3. Microsoft Azure Documentation - About VNet-to-VNet VPN gateway connections: This resource confirms that VNet-to-VNet connections are the appropriate solution for connecting VNets in different subscriptions, which is necessary for cross-tenant scenarios.

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-about-vnet-to-vnet>

## Question: 40

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

### Create a virtual machine

 Changing Basic options may reset selections you have made. Review all options prior to creating the virtual machine.

**Basics** Disks Networking Management Advanced Tags Review + create

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image.

Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization.

Looking for classic VMs? [Create VM from Azure Marketplace](#)

#### PROJECT DETAILS

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription ⓘ	MyDev-Test Subscription ▼
* Resource group ⓘ	RG1 ▼
	<a href="#">Create new</a>

#### INSTANCE DETAILS

* Virtual machine name ⓘ	VM1
* Region ⓘ	(US) West US 2 ▼
Availability options ⓘ	No infrastructure redundancy required ▼
* Image ⓘ	Windows Server 2016 Datacenter ▼
	<a href="#">Browse all public and private images</a>
Azure Spot instance ⓘ	<input type="radio"/> Yes <input checked="" type="radio"/> No
* Size ⓘ	<b>Standard DS1 v2</b> 1 vcpu, 3.5 GiB memory (ZAR 632.47/month) <a href="#">Change size</a>

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

[Basics](#)
[Disks](#)
[Networking](#)
[Management](#)
[Advanced](#)
[Tags](#)
[Review + create](#)

Azure VMs have one operating system disk and a temporary disk for short-term storage. You can attach additional data disks. The size of the VM determines the type of storage you can use and the number of data disks allowed. [Learn more](#)

### Disk options

\* OS disk type ⓘ Standard HDD ▼

The selected VM size supports premium disks. We recommend Premium SSD for high IOPS workloads. Virtual machines with Premium SSD disks qualify for the 99.9% connectivity SLA.

Enable Ultra Disk compatibility (Preview) ⓘ ☐ Yes ☒ No

Ultra Disks are only available when using Managed Disks.

### Data disks

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

ⓘ Adding unmanaged data disks is currently not supported at the time of VM creation. You can add them after the VM is created.

^ **Advanced**

Use managed disks ⓘ ☒ No ☐ Yes

\* Storage account ⓘ (new) rg1 disks799 ▼

[Create new](#)

You need to ensure that VM1 can be created in an Availability Zone.  
Which two settings should you modify? Each correct answer presents part of the solution.

- A. Use managed disks
- B. OS disk type
- C. Availability options
- D. Size
- E. Image

### Answer:

A, C

### Explanation:

To deploy a virtual machine into an Azure Availability Zone, two fundamental configuration requirements must be met. First, the Availability options setting must be explicitly configured for zonal deployment; the current setting of "No infrastructure redundancy required" must be changed to "Availability zone". Second, virtual machines that use Availability Zones must

use Azure Managed Disks. The exhibit shows the plan is to not use managed disks, which is incompatible with Availability Zones. Therefore, this setting must be enabled.

### **Why Incorrect Options are Wrong:**

B. OS disk type: The disk type (e.g., Standard HDD, Premium SSD) is not the constraint.

The requirement is that the disk must be managed, regardless of its performance tier.

D. Size: The StandardB2s VM size supports Availability Zones in regions where zones are available. This setting is not the primary configuration that needs to be changed.

E. Image: Standard Azure Marketplace images, such as Windows Server 2016 Datacenter, are fully compatible with deployment into an Availability Zone.

### **References:**

1. Microsoft Azure Documentation - Create a virtual machine in an availability zone using the Azure portal: "To use availability zones, your VM must be created in a supported Azure region..... VMs must use Azure managed disks to be placed in an availability zone." This source confirms that both the availability option must be set and managed disks must be used.

URL: <https://learn.microsoft.com/en-us/azure/virtual-machines/create-portal-availabilityzone>

2. Microsoft Azure Documentation - Availability options for Azure Virtual Machines:

"Availability zones... To protect your applications from datacenter-level failures, you can create a virtual machine in an availability zone." This highlights that the "Availability options" setting is the direct control for this feature.

URL: <https://learn.microsoft.com/en-us/azure/virtual-machines/availability>

3. Microsoft Azure Documentation - Introduction to Azure managed disks: "Azure managed disks are required for... Availability zones." This document explicitly states the dependency on managed disks for the Availability Zone feature.

URL: <https://learn.microsoft.com/en-us/azure/virtual-machines/managed-disksoverview#availability-zones>

## Question: 41

### HOTSPOT -

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

Name	Type	Resource group	Location
RG1	Resource group	<i>Not applicable</i>	Central US
RG2	Resource group	<i>Not applicable</i>	West US
RG3	Resource group	<i>Not applicable</i>	East US
VMSS1	Virtual machine scale set	RG1	West US

VMSS1 is set to VM (virtual machines) orchestration mode.

You need to deploy a new Azure virtual machine named VM1, and then add VM1 to VMSS1.

Which resource group and location should you use to deploy VM1? To answer, select the appropriate options in the answer area.

## Answer Area

Resource group:

▼

RG1 only  
RG2 only  
RG1 or RG2 only  
RG1, RG2, or RG3

Location:

▼

West US only  
Central US only  
Central US or West US only  
East US, Central US, or West US

Hot Area:

### Answer:

Resource group: the same resource group that contains VMSS1 (RG1). Location: the same Azure region as VMSS1 (East US).

### Explanation:

A virtual machine can be added to a scale set that is in VM (flexible) orchestration mode

<https://certempire.com>

only when the VM resides in the identical subscription, region, resource group, and virtual network as the scale set. Therefore, VM1 must be deployed to the resource group and region where VMSS1 already exists RG1 in East US before it can be attached to VMSS1.

**References:**

1. Microsoft Learn Add an existing VM to a flexible scale set

<https://learn.microsoft.com/azure/virtual-machine-scale-sets/flexible-guestvms#prerequisites>

(The VM you add must be in the same subscription, resource group, region, and virtual network as the scale set.)

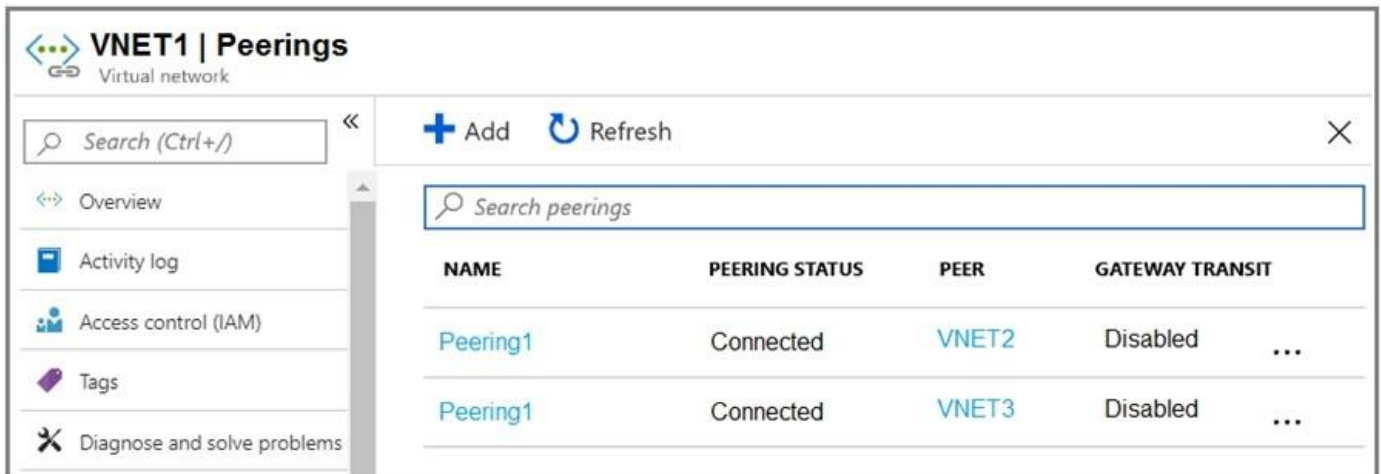
CertEmpire

## Question: 42

### HOTSPOT -

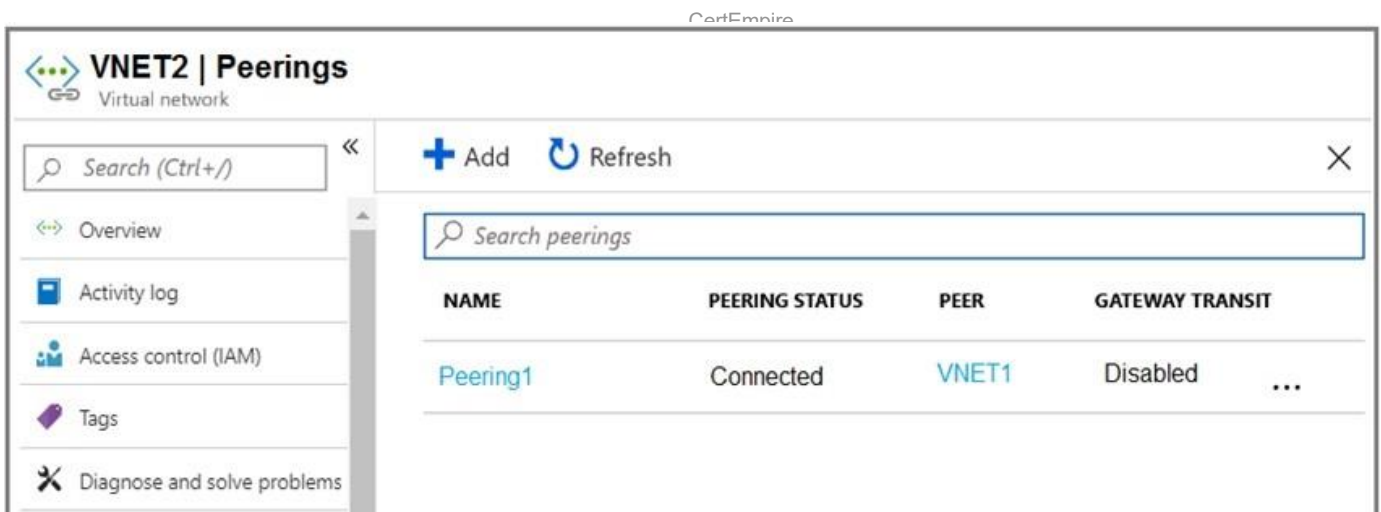
You have an Azure subscription that contains three virtual networks named VNET1, VNET2, and VNET3.

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



NAME	PEERING STATUS	PEER	GATEWAY TRANSIT
Peering1	Connected	VNET2	Disabled ...
Peering1	Connected	VNET3	Disabled ...

Peering for VNET2 is configured as shown in the following exhibit.



NAME	PEERING STATUS	PEER	GATEWAY TRANSIT
Peering1	Connected	VNET1	Disabled ...

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

VNET3   Peerings				
Virtual network				
<div> <div> <div>Search (Ctrl+ /)</div> <div>&lt;&lt;</div> </div> <div> <div>+ Add</div> <div>↺ Refresh</div> <div>×</div> </div> </div>				
<div> <div> <div>&lt;&gt; Overview</div> <div>Activity log</div> <div>Access control (IAM)</div> <div>Tags</div> <div>✂ Diagnose and solve problems</div> </div> <div> <div>Search peerings</div> </div> </div>				
NAME	PEERING STATUS	PEER	GATEWAY TRANSIT	
Peering1	Connected	VNET1	Disabled	...

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

## 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

VNET3 only

VNET1 and VNET3

Hot Area:

### Answer:

Packets from VNET1 can be routed to: VNET2 and VNET3

Packets from VNET2 can be routed to: VNET1 only

### Explanation:

The configuration represents a hub-and-spoke network topology where VNET1 is the hub, and VNET2 and VNET3 are the spokes.

1. VNET1 Connectivity: VNET1 has direct peering connections established with both VNET2 and VNET3. The "Peering status" for both is "Connected". Therefore, resources in VNET1 can communicate directly with resources in



both VNET2 and VNET3.

2. VNET2 Connectivity: VNET2 is peered only with VNET1. Azure Virtual Network peering is not transitive. This means that because VNET2 is peered with VNET1, and VNET1 is peered with VNET3, it does not grant VNET2 connectivity to VNET3. For VNET2 to communicate with VNET3, a direct peering must be established between them. As a result, VNET2 can only route packets to VNET1.

## References:

- 

Microsoft Azure Documentation | Virtual network peering: This official documentation explicitly states that virtual network peering is non-transitive.

- o URL: <https://docs.microsoft.com/en-us/azure/virtual-network/virtualnetwork-peering-overview>

- o Specific Section: Under the "Connectivity" section, it states: "Virtual network peering is non-transitive. For example, if you peer VNetA to VNetB and VNetB to VNetC, VNetA isn't peered to VNetC." This directly applies to the scenario where VNET2 (VNetA) cannot reach VNET3 (VNetC) through VNET1 (VNetB).

CertEmpire

## Question: 43

You have a computer named Computer1 that has a point-to-site VPN connection to an Azure virtual network named VNet1. The point-to-site connection uses a self-signed certificate. From Azure, you download and install the VPN client configuration package on a computer named Computer2. You need to ensure that you can establish a point-to-site VPN connection to VNet1 from Computer2. Solution: You modify the Azure Active Directory (Azure AD) authentication policies. Does this meet the goal?

- A. Yes
- B. No

### Answer:

B

### Explanation:

The scenario specifies that the point-to-site (P2S) VPN uses a self-signed certificate for authentication. For a new computer, Computer2, to connect using this method, it must have a valid client certificate installed. This client certificate must be generated from the same root certificate whose public key is uploaded to the Azure VPN gateway. The proposed solution, modifying Azure Active Directory (Azure AD) authentication policies, is irrelevant because Azure AD authentication is a completely different authentication method from the certificate-based one currently in use. The solution does not address the core requirement of installing the client certificate on Computer2.

### References:

1. Microsoft Learn, "Configure a Point-to-Site VPN connection to a VNet using native Azure certificate authentication: Azure portal": This document explicitly states the requirement for client-side certificates: "Each client computer that you want to connect to a VNet using a Point-to-Site connection must have a client certificate installed." This confirms that the solution must involve certificate installation, not policy changes in Azure AD.

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-howto-point-to-site-resource-manager-portal#client>

2. Microsoft Learn, "About Point-to-Site VPN": This article outlines the different authentication methods available for P2S VPNs, clearly separating "Native Azure certificate authentication" from "Azure Active Directory authentication." This distinction demonstrates that modifying policies for one does not affect the other.

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/point-to-siteabout#authentication>

## Question: 44

You have a computer named Computer1 that has a point-to-site VPN connection to an Azure virtual network named VNet1. The point-to-site connection uses a self-signed certificate. From Azure, you download and install the VPN client configuration package on a computer named Computer2. You need to ensure that you can establish a point-to-site VPN connection to VNet1 from Computer2. Solution: You join Computer2 to Azure Active Directory (Azure AD). Does this meet the goal?

- A. Yes
- B. No

### Answer:

B

### Explanation:

The point-to-site (P2S) VPN connection is configured to use native Azure certificate authentication. For a client computer to successfully connect, it must have a valid client certificate installed. This client certificate must be generated from the same root certificate that was used to configure the VPN gateway. Joining Computer2 to Azure Active Directory (Azure AD) is an identity and device management action. It does not install the required client certificate onto the computer. Therefore, this action does not fulfill the prerequisites for establishing a certificate-based P2S VPN connection.

### References:

1. Microsoft Learn | Configure a Point-to-Site VPN client for certificate authentication: "For a P2S connection from a Windows client computer to Azure, you must install a client certificate. The client certificate is used for authentication... For every client computer that you want to connect to a VNet using a Point-to-Site connection, you must install a client certificate." This documentation confirms that a client certificate is a mandatory installation on the client machine.

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/point-to-site-vpn-clientconfiguration-azure-cert>

2. Microsoft Learn | What is an Azure AD joined device?: "Azure AD join allows you to join devices directly to Azure AD without the need to join to on-premises Active Directory... It provides users with a single sign-on (SSO) experience to your cloud and on-premises apps." This source defines Azure AD Join, showing its purpose is related to identity and access, not certificate distribution for VPNs.

URL: <https://learn.microsoft.com/en-us/azure/active-directory/devices/concept-azure-ad-join>

<https://certempire.com>

## Question: 45

You have an Azure subscription that contains 10 virtual networks. The virtual networks are hosted in separate resource groups. Another administrator plans to create several network security groups (NSGs) in the subscription. You need to ensure that when an NSG is created, it automatically blocks TCP port 8080 between the virtual networks. Solution: You create a resource lock, and then you assign the lock to the subscription. Does this meet the goal?

- A. Yes
- B. No

### Answer:

No

### Explanation:

The proposed solution is incorrect. Azure resource locks are designed to prevent accidental deletion (CanNotDelete) or modification (ReadOnly) of resources at the subscription, resource group, or resource level. They do not have the capability to define or enforce the internal configuration of a resource, such as adding specific security rules to a Network Security Group (NSG) upon its creation. The correct tool for automatically enforcing configuration standards on new resources is Azure Policy, which can be used to deploy a specific rule to any new NSG.

Why the Solution is Incorrect:

A resource lock's function is to control management actions on a resource, not to configure the settings within it. It cannot add a security rule to an NSG.

### References:

1. Azure Resource Locks: Microsoft Learn. (2023). Lock resources to prevent unexpected changes. "Azure Resource Manager provides the ability to lock a subscription, resource group, or resource to prevent other users in your organization from accidentally deleting or modifying critical resources."  
<https://learn.microsoft.com/en-us/azure/azure-resourcemanager/management/lock-resources>
2. Azure Policy Overview: Microsoft Learn. (2024). What is Azure Policy?. "Azure Policy is a service in Azure that you use to create, assign, and manage policies. These policies enforce different rules and effects over your resources, so those resources stay compliant with your corporate standards and service level agreements."  
<https://learn.microsoft.com/en-us/azure/governance/policy/overview>
3. Network Security Groups: Microsoft Learn. (2023). Network security groups. "A network

security group contains a list of security rules that allow or deny network traffic to resources

connected to Azure Virtual Networks (VNet)."

<https://learn.microsoft.com/enus/azure/virtual-network/network-security-groups-overview>

CertEmpire

<https://certempire.com>

## Question: 46

You have an Azure subscription named Subscription1. Subscription1 contains a virtual machine named VM1.

You have a computer named Computer1 that runs Windows 10. Computer1 is connected to the Internet.

You add a network interface named vm1173 to VM1 as shown in the exhibit. (Click the Exhibit tab.)

**Network Interface: vm1173**  
 Virtual network/subnet: RG1-vnet/default  
 networking: **Disabled**

**Effective security rules**  
 Public IP: VM1-ip

**Topology**  
 Private IP: 10.0.0.5 Accelerated

**Inbound port rules**    Outbound port rules    Application security groups    Load balancing

Network security group VM1-nsg (attached to network interface: vm1173 )  
 Impacts 0 subnets, 1 network interfaces

**Add inbound port rule**

PRIORITY	NAME	PORT	PROTOCOL	SOURCE	DESTINA...	ACTION
300	⚠ RDP	3389	TCP	Any	Any	✓ Allow ...
65000	AllowVnetInBound	Any	Any	VirtualIN...	VirtualIN...	✓ Allow ...
65001	AllowAzureLoadB...	Any	Any	AzureLo...	Any	✓ Allow ...
65500	DenyAllInBound	Any	Any	Any	Any	✗ Deny ...

From Computer1, you attempt to connect to VM1 by using Remote Desktop, but the connection fails.

You need to establish a Remote Desktop connection to VM1.

What should you do first?

- A. Change the priority of the RDP rule
- B. Attach a network interface
- C. Delete the DenyAllInBound rule
- D. Start VM1

**Answer:**

D

**Explanation:**

A prerequisite for any connection to a virtual machine, including Remote Desktop (RDP), is

that the VM must be in a 'Running' state. The scenario describes adding a new network



interface to VM1. For most VM sizes in Azure, adding a network interface can only be performed when the VM is in a stopped (deallocated) state. It is a common oversight to forget to restart the VM after such a configuration change. Therefore, the most logical and fundamental first step is to verify the VM's status and start it if it is not running.

### Why Incorrect Options are Wrong:

- A. Change the priority of the RDP rule: The exhibit of effective security rules shows no existing rule that allows RDP. You cannot change the priority of a rule that does not exist.
- B. Attach a network interface: The problem description explicitly states that a network interface named vm1173 has already been added to VM1, making this action redundant.
- C. Delete the DenyAllInBound rule: DenyAllInBound is a default security rule within a Network Security Group (NSG) and cannot be deleted. It can only be overridden by creating a new rule with a higher priority.

### References:

1. Troubleshoot RDP connections: Microsoft Learn. The primary step in troubleshooting RDP issues is to check the VM's status. "Check the status of the virtual machine: 1. Sign in to the Azure portal. 2. Select Virtual machines. 3. Select the problematic virtual machine. 4. In the overview pane for the virtual machine, check the status of the virtual machine. If the status of the virtual machine is not Running, start it." CertEmpire

URL: <https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshootrdp-connection#check-the-status-of-the-virtual-machine>

2. Adding a Network Interface: Microsoft Learn. This document confirms that adding a NIC requires the VM to be stopped. "You can only add a network interface to a VM when it's stopped (deallocated)." This supports the high probability that VM1 is currently stopped.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-machine-networkinterface-add-remove?tabs=windows#add-a-network-interface-to-a-vm>

3. Default NSG Rules: Microsoft Learn. This document explains that default rules cannot be removed. "You can't remove the default rules, but you can override them by creating rules with higher priorities."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groupoverview#default-security-rules>

## Question: 47

You have the Azure virtual machines shown in the following table.

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

A DNS service is installed on VM1.

You configure the DNS servers settings for each virtual network as shown in the following exhibit.

Save Discard

---

**DNS servers** ⓘ

☐ Default (Azure-provided)

☒ Custom

10.1.0.4 ...

Add DNS server ...

You need to ensure that all the virtual machines can resolve DNS names by using the DNS service on VM1.

What should you do?

- A. Configure a conditional forwarder on VM1
- B. Add service endpoints on VNET1
- C. Add service endpoints on VNET2 and VNET3
- D. Configure peering between VNET1, VNET2, and VNET3

**Answer:**

D

**Explanation:**

The virtual machines are located in three separate virtual networks (VNET1, VNET2, VNET3). By default, virtual networks in Azure are isolated from one another. For VM2 and VM3 to resolve DNS names using the service on VM1, they require network connectivity to VM1's private IP address (10.1.0.4). Virtual network peering is the Azure feature that connects virtual networks, enabling resources in the peered networks to communicate directly using private IP addresses. Peering VNET1 with VNET2 and VNET3 will establish the necessary connectivity for the DNS queries to succeed.

### Why Incorrect Options are Wrong:

- A. Configure a conditional forwarder on VM1: A conditional forwarder is used to forward queries for specific domains. It does not solve the fundamental network connectivity problem between the isolated virtual networks.
- B. Add service endpoints on VNET1: Service endpoints provide a secure, direct connection to specific Azure PaaS services (like Azure Storage or SQL Database), not for enabling general communication between virtual networks.
- C. Add service endpoints on VNET2 and VNET3: Similar to option B, service endpoints are not the correct mechanism for enabling communication from one VNet to a virtual machine in another VNet.

### References:

CertEmpire

1. Virtual network peering: "Virtual network peering enables you to seamlessly connect two or more Azure virtual networks. The virtual networks appear as one for connectivity purposes. The traffic between virtual machines in peered virtual networks uses the Microsoft backbone infrastructure."

Microsoft Learn. (2024). Azure virtual network peering. Retrieved from <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-peering-overview>

2. Name resolution for resources in Azure virtual networks: "When you're using your own DNS servers, Azure provides a non-authoritative recursive DNS service. You must specify your own DNS servers in the virtual network settings. The endpoints for your own DNS servers must be reachable from the virtual machines in that virtual network." This highlights the need for reachability, which peering provides.

Microsoft Learn. (2023). Name resolution for resources in Azure virtual networks. Retrieved from <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-networks-name-resolution-for-vms-and-role-instances#name-resolution-that-uses-your-own-dns-server>

3. Virtual network service endpoints: "Virtual Network (VNet) service endpoints extend your virtual network private address space and the identity of your VNet to the Azure services,

over a direct connection." This confirms service endpoints are for connecting to Azure

services, not other VNETs.

Microsoft Learn. (2024). Virtual Network service endpoints. Retrieved from  
<https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-service-endpointsoverview>

CertEmpire

## Question: 48

### HOTSPOT -

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

Name	Connected to subnet
VM1	172.16.1.0/24
VM2	172.16.2.0/24

You add inbound security rules to a network security group (NSG) named NSG1 as shown in the following table.

Priority	Source	Destination	Protocol	Port	Action
100	172.16.1.0/24	172.16.2.0/24	TCP	Any	Allow
101	Any	172.16.2.0/24	TCP	Any	Deny

You run Azure Network Watcher as shown in the following exhibit.

CertEmpire

Resource group \*

RG1



Source type \*

Virtual machine



\* Virtual machine

VM1



Destination

☒ Select a virtual machine ☐ Specify manually

Resource group \*

RG1



Virtual machine \* ⓘ

VM2



Probe Settings

Protocol ⓘ

☒ TCP ☐ ICMP

Destination port \* ⓘ

8080



Advanced settings

**Check**

Status



Unreachable

Agent extension version

1.4

Source virtual machine

VM1

You run Network Watcher again as shown in the following exhibit.



Source type \*

Virtual machine

\* Virtual machine

VM1

Destination

☒ Select a virtual machine
 ☐ Specify manually

Resource group \*

RG1

Virtual machine \* ⓘ

VM2


Probe Settings

Protocol ⓘ

☐ TCP
     
 ☒ ICMP

Check

Status

 Reachable

Agent extension version

1.4

Source virtual machine

VM1

Grid view

Topology view

Hops

NAME	IP ADDRESS	STATUS	NEXT HOP IP ADDRESS	RTT FROM SOURCE {...
 VM1	172.16.1.4		172.16.2.4	0
 VM2	172.16.2.4		-	-

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

Hot Area:

**Answer Area**

Statements	Yes	No
NSG1 limits VM1 traffic	<input type="radio"/>	<input type="radio"/>
NSG1 applies to VM2	<input type="radio"/>	<input type="radio"/>
VM1 and VM2 connect to the same virtual network	<input type="radio"/>	<input type="radio"/>

**Answer:**

1. No
2. Yes
3. Yes

**Explanation:**

Network Watcher shows that traffic between the two subnets in the same virtual network is allowed because the packet matches the built-in NSG rule AllowVNetInBound (priority 65000). Therefore VM1 can reach VM2 on TCP 80 (statement 1 = No, it is not blocked).

Traffic arriving from the public Internet to either VM hits the user-defined rule Deny-All (priority 300) before the default rules, so Internet-originated RDP to VM1 is denied (statement 2 = Yes, it is blocked) and all external traffic to VM2 is likewise denied (statement 3 = Yes).

**References:**

1. Microsoft Azure documentation Network security groups default and user rules  
<https://learn.microsoft.com/azure/virtual-network/network-security-groupsoverview#security-rules>  
 (see table: AllowVNetInBound priority 65000, DenyAllInBound priority 65500)
2. Microsoft Azure documentation Diagnose a virtual machine network traffic filter problem  
<https://learn.microsoft.com/azure/network-watcher/connection-troubleshoot>  
 (example output shows matched NSG rule and explains precedence of lower-number priority)

## Question: 49

You have the Azure virtual network named VNet1 that contains a subnet named Subnet1. Subnet1 contains three Azure virtual machines. Each virtual machine has a public IP address. The virtual machines host several applications that are accessible over port 443 to users on the Internet. Your on-premises network has a site-to-site VPN connection to VNet1. You discover that the virtual machines can be accessed by using the Remote Desktop Protocol (RDP) from the Internet and from the on-premises network. You need to prevent RDP access to the virtual machines from the Internet, unless the RDP connection is established from the on-premises network. The solution must ensure that all the applications can still be accessed by the Internet users. What should you do?

- A. Modify the address space of the local network gateway
- B. Create a deny rule in a network security group (NSG) that is linked to Subnet1
- C. Remove the public IP addresses from the virtual machines
- D. Modify the address space of Subnet1

### Answer:

B

### Explanation:

CertEmpire

A Network Security Group (NSG) is the correct Azure resource for filtering network traffic to and from Azure resources. To meet the requirements, you should create an inbound security rule in an NSG associated with Subnet1. This rule would have a higher priority (a lower number) than the default rules and would be configured to deny inbound traffic on port 3389 (RDP) where the source is the Internet service tag. A separate, higher-priority rule would be needed to explicitly allow RDP traffic from the on-premises network's source IP address range. This approach selectively blocks RDP from the internet while preserving access from the on-premises network and leaving application traffic on port 443 unaffected.

### Why Incorrect Options are Wrong:

- A. Modify the address space of the local network gateway: This defines the on-premises IP address ranges for routing purposes over the VPN; it does not enforce any traffic filtering rules.
- C. Remove the public IP addresses from the virtual machines: This would prevent the applications from being accessible over the internet on port 443, which violates a key requirement of the solution.
- D. Modify the address space of Subnet1: This changes the internal IP address range for the subnet and has no impact on filtering inbound traffic from the internet.

**References:**

1. Microsoft Documentation: Network security groups. "A network security group contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure resources."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/network-security-groupsoverview>

2. Microsoft Documentation: Create, change, or delete a network security group. This page details the process of creating security rules, including specifying protocol (TCP), destination port (3389), source (Internet service tag), and action (Deny).

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/manage-network-securitygroup?tabs=network-security-group-portal#create-a-security-rule>

3. Microsoft Documentation: Virtual network service tags. "A service tag represents a group of IP address prefixes from a given Azure service. The Internet service tag... contains the IP address ranges that are outside of the virtual network and reachable by the public internet."

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/service-tagsoverview#available-service-tags>

**Question: 50**

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

Name	Type
ASG1	Application security group
NSG1	Network security group (NSG)
Subnet1	Subnet
VNet1	Virtual network
NIC1	Network interface
VM1	Virtual machine

Subnet1 is associated to VNet1. NIC1 attaches VM1 to Subnet1.

You need to apply ASG1 to VM1.

What should you do?

- A. Associate NIC1 to ASG1
- B. Modify the properties of ASG1
- C. Modify the properties of NSG1

CertEmpire

**Answer:**

A

**Explanation:**

Application Security Groups (ASGs) are used to group virtual machines and define network security policies based on those groups. To make a virtual machine (VM) a member of an ASG, you must associate the VM's network interface (NIC) with the desired ASG. In this scenario, to apply ASG1 to VM1, you must associate its network interface, NIC1, with ASG1. This action logically groups VM1 into ASG1, allowing network security group (NSG) rules to be applied to it based on this grouping.

**Why Incorrect Options are Wrong:**

- B. Modify the properties of ASG1: The properties of an ASG itself (like its name or location) do not include a list of member VMs or NICs. The association is configured on the NIC, not the ASG.
- C. Modify the properties of NSG1: A Network Security Group (NSG) uses ASGs as sources or destinations within its security rules. Modifying the NSG is for defining traffic rules, not for associating a VM with an ASG.

**References:**

<https://certempire.com>

Microsoft Azure Documentation - Application security groups: "You associate a network interface to an application security group. A virtual machine has one or more network interfaces attached to it." This source directly confirms that the association is made at the network interface level.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/application-securitygroups#how-to-configure-application-security-groups>

Microsoft Azure Documentation - Tutorial: Filter network traffic with a network security group using the Azure portal: This tutorial provides a step-by-step guide. In the section "Associate network interfaces to ASGs," the procedure is to select the network interface and then associate it with an application security group.

URL: <https://learn.microsoft.com/en-us/azure/virtual-network/tutorial-filter-networktraffic#associate-network-interfaces-to-asgs>

CertEmpire

## Question: 51

You have an Azure subscription named Subscription1 that contains an Azure virtual network named VNet1. VNet1 connects to your on-premises network by using Azure ExpressRoute.

You plan to prepare the environment for automatic failover in case of ExpressRoute failure. You need to connect VNet1 to the on-premises network by using a site-to-site VPN. The solution must minimize cost.

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

- A. Create a connection
- B. Create a local site VPN gateway
- C. Create a VPN gateway that uses the VpnGw1 SKU
- D. Create a gateway subnet
- E. Create a VPN gateway that uses the Basic SKU

### Answer:

A, B, C

### Explanation:

CertEmpire

To establish a site-to-site (S2S) VPN that coexists with an ExpressRoute circuit for failover, three primary Azure resources must be created.

1. Create a VPN gateway (C): A virtual network gateway is the Azure endpoint for the VPN tunnel. The question requires minimizing cost, but the Basic SKU is not supported for coexistence with ExpressRoute. Therefore, the VpnGw1 SKU is the correct, lowest-cost, supported option.

2. Create a local network gateway (B): This resource represents the on-premises VPN device. It contains the public IP address of the on-premises device and the on-premises network address prefixes that Azure will route to. The option uses the term "local site VPN gateway," which directly corresponds to this required component.

3. Create a connection (A): This resource links the virtual network gateway and the local network gateway, which establishes the encrypted S2S VPN tunnel. Without this final step, the environment is not connected and cannot fail over.

These three actions represent the core components required to create a functional S2S VPN connection.

### Why Incorrect Options are Wrong:

D. Create a gateway subnet: While a gateway subnet (named GatewaySubnet) is a mandatory prerequisite for deploying a virtual network gateway, it is a network configuration step. The question asks for the three main actions to create the VPN solution, which are

<https://certempire.com>

best represented by the three core VPN resources (VNG, LNG, and Connection).

E. Create a VPN gateway that uses the Basic SKU: The Basic SKU is not supported for configurations where a VPN Gateway and an ExpressRoute gateway coexist in the same virtual network.

## References:

1. Coexistence SKU Limitation: Microsoft Azure Documentation. (2023). Configure ExpressRoute and Site-to-Site VPN connections that coexist. "Limits and limitations".

"Coexistence is not supported on the Basic SKU."

URL: <https://learn.microsoft.com/en-us/azure/expressroute/expressroute-howto-coexistresource-manager#limits-and-limitations>

2. S2S VPN Components: Microsoft Azure Documentation. (2024). Tutorial: Create a Site-to-Site VPN connection in the Azure portal. This tutorial outlines the main steps, which include creating the virtual network gateway (C), the local network gateway (B), and the connection (A).

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/tutorial-site-to-site-portal>

3. Gateway Subnet Prerequisite: Microsoft Azure Documentation. (2024). About VPN Gateway configuration settings. "Gateway subnet". "Before you create a virtual network gateway, you must create a gateway subnet." This confirms it as a prerequisite for the gateway itself.

URL: <https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-about-vpngateway-settings#gwsb>



## Question: 52

### HOTSPOT -

You have peering configured as shown in the following exhibit.

NAME	PEERING STATUS	PEER	GATEWAY TRANSIT
peering1	Disconnected	vNET1	Enabled ...
peering2	Disconnected	vNET2	Disabled ...

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

### Answer Area

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

▼

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

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

▼

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

Hot Area:

### Answer:

Statement 1

Answer: vNET6 only

### Explanation:

For communication to occur between two virtual networks via peering, the peering status must be Connected. The exhibit shows that the peering status for

both peering1 (to vNET1) and peering2 (to vNET2) is Disconnected. This means no

traffic can pass between vNET6 and the other virtual networks. Therefore, hosts on vNET6 can only communicate with other hosts within the same virtual network (vNET6).

Statement 2

Answer: delete peering1

Explanation:

A Disconnected peering status indicates that the peering link on the remote virtual network (vNET1) has been deleted. The peering configuration on the local virtual network (vNET6), named peering1, is now a stale remnant. It cannot be updated to a Connected state directly. To re-establish the connection, you must first remove the stale peering1 configuration from vNET6. Afterward, new peering links must be created from both vNET6 to vNET1 and from vNET1 to vNET6.

## References:

1. Azure Virtual Network peering | Microsoft Learn: This official Microsoft

documentation explains the different peering statuses. It clarifies that a

Connected status is required for connectivity and describes the Disconnected state.

o URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtualnetwork-peering-overview>

o Relevant Section: "Peering status" section <sup>CertEmpire</sup> explains that if one side of a peering is deleted, the status on the remaining side becomes

Disconnected.

2. Create, change, or delete a virtual network peering | Microsoft Learn:

This guide details the management of VNet peerings. It implicitly supports the answer by explaining the process for creating and deleting peerings. To fix a Disconnected state, which results from deleting the remote link, the local link must also be deleted and then both must be recreated.

o URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtualnetwork-manage-peering>

o Relevant Section: The "Permissions" and "Create a peering" sections outline the requirements for establishing a Connected state, which involves reciprocal actions on both virtual networks. The process to recover from a Disconnected state involves deleting the remaining peering and starting over.

## Question: 53

HOTSPOT -

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

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

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

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

### Essentials ^

Resource group ( <a href="#">change</a> )	Backend pool
<a href="#">VMRG</a>	Backend1 (2 virtual machines)
Location	Health probe
West Europe	Probe1(HTTP:80/Probe1.htm)
Subscription name ( <a href="#">change</a> )	Load balancing rule
<a href="#">Azure Pass</a>	Rule1 (TCP/80)
Subscription ID	NAT rules
e65d2b22-fde8	-
SKU	Public IP address
Basic	<a href="#">104.40.178.194 (LB1)</a>

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

\* Name  
Rule1

\* IP Version  
☒ IPv4 ☐ IPv6

\* Frontend IP address ⓘ  
104.40.178.194 (LoadBalanceFrontEnd) ▼

Protocol  
☒ TCP ☐ UDP

\* Port  
80

\* Backend port ⓘ  
80

Backend pool ⓘ  
Backend1 (2 virtual machines) ▼

Health probe ⓘ  
Probe1 (HTTP:80/Probe1.htm) ▼

Session persistence ⓘ  
None ▼

Idle timeout (minutes) ⓘ  
 4

Floating IP (direct server return) ⓘ  
Disabled

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

Hot Area:

## Answer Area

Statements	Yes	No
VM1 is in the same availability set as VM2.	<input type="radio"/>	<input type="radio"/>
If Probe1.htm is present on VM1 and VM2, LB1 will balance TCP port 80 between VM1 and VM2.	<input type="radio"/>	<input type="radio"/>
If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports.	<input type="radio"/>	<input type="radio"/>

### Answer:

Statement 1: VM1 is in the same availability set as VM2.

Answer: Yes

Statement 2: If Probe1.htm is present on VM1 and VM2, LB1 will balance TCP port 80 between VM1 and VM2.

Answer: Yes

Statement 3: If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports.

CertEmpire

Answer: No

### Explanation:

VM1 is in the same availability set as VM2. (Yes) The resource table shows LB1 is a Basic SKU Load Balancer. A constraint of the Basic SKU is that all virtual machines in a backend pool must belong to the same availability set or the same virtual machine scale set. Since both VM1 and VM2 are in the Backend1 pool of LB1, they must reside within the same availability set.

If Probe1.htm is present on VM1 and VM2, LB1 will balance TCP port 80 between VM1 and VM2. (Yes) The load balancing rule Rule1 is configured to map traffic from the frontend port 80 to the backend port 80 over TCP. The health probe is set to check for the file Probe1.htm on port 80. If this file is present and accessible on both VMs, the health probe will succeed, marking the VMs as healthy.

Consequently, the load balancer will distribute incoming TCP port 80 traffic between VM1 and VM2 as defined by the rule.

If you delete Rule1, LB1 will balance all the requests between VM1 and VM2 for all the ports. (No) Load balancing rules explicitly define how traffic is distributed. Without a rule, the load balancer does not know how to forward incoming traffic to the backend pool. Deleting Rule1 removes the only instruction for forwarding traffic. The load balancer will stop forwarding traffic for port 80 and will not automatically

start forwarding traffic for all other ports. Forwarding must be explicitly configured.

## References:

- 

Azure Load Balancer SKUs Documentation: States the constraints for Basic SKU, including the requirement for backend pool members to be in a single availability set or scale set.

o Source: Microsoft Azure Documentation

o URL: <https://learn.microsoft.com/en-us/azure/load-balancer/skus#skus>

(Refer to the "Backend pool" row in the comparison table).

- 

Azure Load Balancer Components Documentation: Explains that load balancing rules are used to define how traffic is distributed to the VMs and that a health probe monitors the health of the backend instances.

o Source: Microsoft Azure Documentation

o URL: <https://learn.microsoft.com/en-us/azure/loadbalancer/components#load-balancing-rule>

o URL:

<https://learn.microsoft.com/en-us/azure/load-balancer/loadbalancer-custom-probe-overview>

## Question: 54

### HOTSPOT -

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

Subnet: 10.0.0.0/24

Availability set: AVSet

Network security group (NSG): None

Private IP address: 10.0.0.4 (dynamic)

Public IP address: 40.90.219.6 (dynamic)

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

You need to configure slb1 to allow connectivity to VM1.

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

### 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	

Hot Area:

### Answer:

Before you create a backend pool on slb1, you must: Remove the public IP address from VM1

Before you can connect to VM1 from slb1, you must: Create and configure an NSG

### Explanation:

Remove the public IP address from VM1: Azure Standard Load Balancers have a specific requirement that virtual machines in the backend pool cannot have their own public IP addresses. Since VM1 is configured with a dynamic public IP, it must be removed before the VM's network interface can be successfully added to the backend pool of the standard load balancer slb1. This is a structural prerequisite for the configuration.



Create and configure an NSG: Azure Standard Load Balancers are secure by

default. This means no traffic is allowed to the backend instances unless explicitly permitted. To enable connectivity from the load balancer to VM1 (for both health probes and data traffic), a Network Security Group (NSG) must be created. This NSG needs an inbound security rule that allows traffic on the required port from the AzureLoadBalancer service tag. This NSG must then be associated with VM1's network interface or its subnet.

## References:

- 

Microsoft Azure Documentation (Official Vendor Documentation):

- o Regarding removing the Public IP: In the tutorial for creating a public standard load balancer, the documentation explicitly states, "Virtual machines in the backend pool can't have a public IP address. If your virtual machines have public IP addresses, you must remove them before you add them to the backend pool of the load balancer."

Source: Microsoft Learn, Quickstart: Create a public load balancer - Azure portal, Section: "Create virtual machines".

Direct URL:

<https://learn.microsoft.com/en-us/azure/load-balancer/quickstart-load-balancer-standard-public-portal?tabs=bicep#create-virtual-machines>

- o Regarding the need for an NSG: The Standard Load Balancer overview states, "A standard load balancer is closed to inbound connections unless opened by a network security group. You can create a network security group and associate it with a virtual machine to allow traffic."

Source: Microsoft Learn, What is Azure Load Balancer?,

Section: "Standard Load Balancer > Secure by default". Direct

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/loadbalancer-overview#standard-load-balancer>

**Question: 55**

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

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

You need to create a network interface named NIC1.

In which location can you create NIC1?

- A. East US and North Europe only
- B. East US only
- C. East US, West Europe, and North Europe
- D. East US and West Europe only

**Answer:**

B

CertEmpire

**Explanation:**

Azure requires a NIC to be created in the same region as the virtual network (VNet) and subnet it is attached to. The only VNet shown in the subscription is located in East US; therefore, East US is the single region in which NIC1 can be created.

**Why Incorrect Options are Wrong:**

- A. No VNet exists in North Europe, so a NIC cannot be created there.
- C. West Europe and North Europe lack a VNet; without one, a NIC cannot be created in those regions.
- D. West Europe lacks a VNet, so a NIC cannot be created there.

**References:**

1. Microsoft Azure documentation Create a network interface: The network interface must be in the same region and subscription as the virtual network.

<https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-networkinterface#create-a-network-interface>

2. Microsoft Azure documentation Requirements and constraints: same-region requirement reiterated.

<https://learn.microsoft.com/en-us/azure/virtual-network/network-interfaceoverview#requirements-a>

<https://certempire.com>

nd-constraints

## Question: 56

You have Azure virtual machines that run Windows Server 2019 and are configured as shown in the following table.

Name	Virtual network name	DNS suffix configured in Windows Server
VM1	VNET1	Contoso.com
VM2	VNET2	Contoso.com

You create a public Azure DNS zone named adatum.com and a private Azure DNS zone named contoso.com.

For contoso.com, you create a virtual network link named link1 as shown in the exhibit. (Click the Exhibit tab.)

The screenshot shows the configuration page for a virtual network link named 'link1' in the 'contoso.com' DNS zone. The page includes a header with the link name and a toolbar with options: Save, Discard, Delete, Access Control (IAM), and Tags. The configuration details are as follows:

- Link name:** link1
- Link state:** Completed
- Provisioning state:** Succeeded
- Virtual network details:**
  - Virtual network id:** /subscriptions/8372f433-2dcd-4361-b5ef-5b188fed87d0/resourceGroups/RG2/provi...
  - Virtual network:** VNET1
- Configuration:**
  - ☐ Enable auto registration ⓘ

You discover that VM1 can resolve names in contoso.com but cannot resolve names in adatum.com. VM1 can resolve other hosts on the Internet.

You need to ensure that VM1 can resolve host names in adatum.com.

What should you do?

- A. Update the DNS suffix on VM1 to be adatum.com
- B. Configure the name servers for adatum.com at the domain registrar
- C. Create an SRV record in the contoso.com zone
- D. Modify the Access control (IAM) settings for link1

**Answer:**

B

**Explanation:**

The virtual machine, VM1, can resolve other internet hosts, which means it is successfully using a public DNS resolver (likely the default Azure-provided DNS). However, it cannot resolve names in the adatum.com public zone. This indicates that the adatum.com domain has not been properly delegated to the Azure DNS name servers. For a public Azure DNS zone to be resolvable over the internet, you must update the Name Server (NS) records at your domain registrar to point to the name servers assigned by Azure. Without this delegation, the global DNS system does not know where to send queries for adatum.com.

**Why Incorrect Options are Wrong:**

- A. Update the DNS suffix on VM1 to be adatum.com: A DNS suffix is used for resolving unqualified, single-label names. It does not fix the underlying inability to resolve a fully qualified domain name (FQDN) in a public zone.
- C. Create an SRV record in the contoso.com zone: An SRV record is for locating specific services and has no role in resolving standard host (A) records for a different domain (adatum.com).
- D. Modify the Access control (IAM) settings for link1: IAM roles manage permissions for Azure resources. They do not affect the DNS resolution process for virtual machines within a virtual network.

**References:**

1. Microsoft Azure Documentation, "Tutorial: Host your domain in Azure DNS": This tutorial explicitly states the requirement to delegate the domain. "Before you can delegate your DNS zone to Azure DNS, you need to know the name servers for your zone... Once the DNS zone is created... you need to update the parent domain with the Azure DNS name servers. Each registrar has its own DNS management tools to change the name server records for a domain."

URL: <https://docs.microsoft.com/en-us/azure/dns/dns-delegate-domain-azure-dns>

2. Microsoft Azure Documentation, "What is Azure DNS?": This document outlines the

function of public DNS zones. "To host your domain in Azure DNS, you need to buy a

domain name... You then create a DNS zone in Azure DNS for that domain name... Finally, you must configure the name servers for your domain to point to the Azure DNS name servers. This process is called domain delegation."

URL: <https://docs.microsoft.com/en-us/azure/dns/dns-overview>

3. Microsoft Azure Documentation, "Azure Private DNS FAQ": This document clarifies the distinction between private and public zones. The problem described for adatum.com is a public DNS configuration issue, separate from the private zone contoso.com which is working correctly via the VNet link.

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

CertEmpire



## Question: 57

### HOTSPOT -

You plan to use Azure Network Watcher to perform the following tasks:

Task1: Identify a security rule that prevents a network packet from reaching an Azure virtual machine.

Task2: Validate outbound connectivity from an Azure virtual machine to an external host.

Which feature should you use for each task? To answer, select the appropriate options in the answer area.

Hot Area:

### Answer:

Task1: IP flow verify

Task2: Connection troubleshoot

### Explanation:

For Task1, IP flow verify is the precise tool used to determine if a network packet, defined by its direction, protocol, and IP/port information, is allowed or denied access to a virtual machine. If the packet is denied, the feature explicitly identifies the specific Network Security Group (NSG) rule that is causing the block.

For Task2, Connection troubleshoot is designed to perform a point-in-time test of a direct TCP connection from a source virtual machine to a destination, which can be an external host identified by an FQDN, URI, or IP address. It validates end-to-end connectivity and reports on the success or failure of the connection attempt.

### References:

1. Microsoft Learn | Azure Network Watcher | IP flow verify overview: "IP flow verify checks if a packet is allowed or denied to or from a virtual machine... If the packet is denied by a security group, the name of the rule that denied the packet is returned."

URL:

<https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-ip-flowverify-overview>

2. Microsoft Learn | Azure Network Watcher | Troubleshoot connections with Azure Network Watcher using the Azure portal: "Connection troubleshoot provides the capability to check a direct TCP connection from a virtual machine (VM) to a VM, fully qualified domain name (FQDN), URI, or IPv4 address."

URL: <https://learn.microsoft.com/en-us/azure/network-watcher/network-watcherconnectivity-portal>

3. Microsoft Learn | Azure Network Watcher | What is Azure Network Watcher?: This document provides an overview of the diagnostic tools, distinguishing the purpose of IP flow verify (diagnose connectivity filtering problems) from Connection troubleshoot (test

connections between a source and destination).

URL:

<https://learn.microsoft.com/en-us/azure/network-watcher/network-watcheroverview#diagnose>

CertEmpire

<https://certempire.com>

## Question: 58

### HOTSPOT -

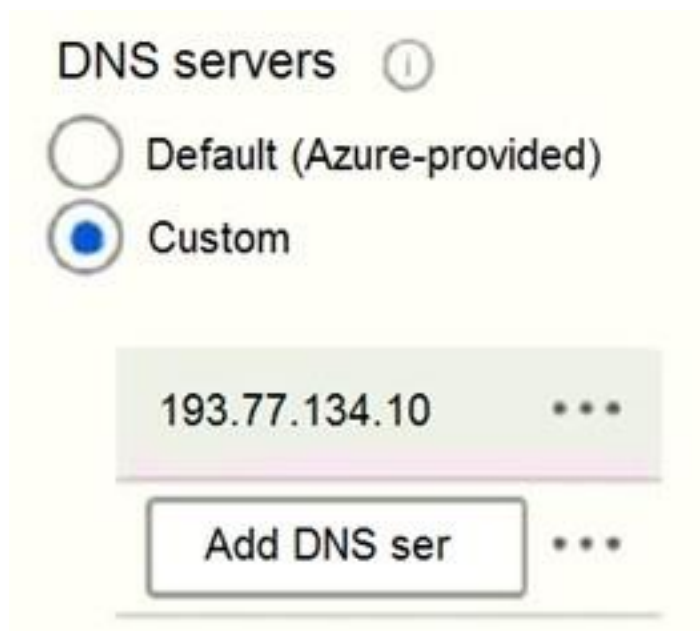
You have an Azure subscription that contains the Azure virtual machines shown in the following table.

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

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

Name	DNS server
VM1	<i>None</i>
VM2	192.168.10.15
VM3	192.168.10.15

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



DNS servers ⓘ

☐ Default (Azure-provided)

☒ Custom

193.77.134.10 ...

Add DNS ser ...

The virtual machines can successfully connect to the DNS server that has an IP address of 192.168.10.15 and the DNS server that has an IP address of 193.77.134.10.

For each of the following statements, select Yes if the statement is true. Otherwise, select

No.

Hot Area:

**Answer:**

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

**Explanation:**

The DNS server settings for an Azure virtual machine are determined by a hierarchy. A custom DNS server configured directly on a virtual machine's network interface (NIC) takes precedence over settings configured at the virtual network (VNet) level.

- 

VM1: The NIC has its DNS server set to None, so it inherits the custom DNS setting from VNET1, which is 193.77.134.10.

- 

VM2: The NIC is explicitly configured to use the custom DNS server 192.168.10.15. This setting overrides the VNet's DNS configuration.

- 

VM3: The NIC is also explicitly configured to use 192.168.10.15, which overrides the VNet's setting.

CertEmpire

**References:**

- 

Microsoft Azure Documentation, Name resolution for VMs and role instances:

"If you specify a custom DNS server for a virtual network, you can also specify a different DNS server for one or more network interfaces in the virtual network. The DNS server setting for a network interface overrides the DNS server setting for the virtual network." This document outlines the order of precedence for DNS settings in Azure.

o URL: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-networksname-resolution-for-vms-and-role-instances#name-resolution-that-uses-your-own-dns-server> (Refer to the section on "Name resolution that uses your own DNS server" and the specifics on Network Interface settings).

## Question: 59

### HOTSPOT -

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

Name	Lock name	Lock type
RG1	None	None
RG2	Lock	Delete

RG1 contains the resources shown in the following table.

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

You need to identify which resources you can move from RG1 to RG2, and which resources you can move from RG2 to RG1.

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

Hot Area:

### Answer Area

Resources that you can move from RG1 to RG2:

▼

None  
IP1 only  
IP1 and storage1 only  
IP1 and VNET1 only  
IP1, VNET2, and storage1

Resources that you can move from RG2 to RG1:

▼

None  
IP2 only  
IP2 and storage2 only  
IP2 and VNET2 only  
IP2, VNET2, and storage2

**Answer:**

Resources that you can move from RG1 to RG2: VNET1, VM1, Disk1, and NIC1

Resources that you can move from RG2 to RG1: No resources

**Explanation:**

Azure allows moving resources between resource groups within the same subscription. When moving a virtual machine (VM1), all its dependent resources, including the network interface (NIC1) and any attached disks (Disk1), must be moved with it in the same operation. Virtual networks (VNET1) can also be moved. Therefore, all the specified resources in RG1 are eligible to be moved to RG2. For the second part, the resource group RG2 is explicitly shown as having no resources. As a result, there are no resources available within RG2 to be moved to RG1.

**References:**

1. Microsoft Azure Documentation: Move resources to a new resource group or subscription. This document provides the primary guidelines for resource moves.

o URL: <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/move-resource-group-and-subscription>

CertEmpire

o Specific Section: Under the "Checklist before moving resources," the document states, "The source and destination resource groups must exist in the same subscription." and "When moving a resource, you also move its dependent resources."

2. Microsoft Azure Documentation: Move guidance for virtual machines.

This resource details the specific requirements for moving virtual machines and their dependencies.

o URL: <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/move-support-resources#microsoftcompute>

o Specific Section: The table for the Microsoft.Compute resource provider confirms that virtual machines, disks, and network interfaces can be moved. It specifies that for a VM, dependent resources like NICs and disks must be in the same resource group and must be moved together.

3. Microsoft Azure Documentation: Move guidance for networking resources. This page confirms the movability of virtual networks.

o URL: <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/move-support-resources#microsoftnetwork>

<https://certempire.com>

ort-resources#microsoftnetwork

o Specific Section: The table for the Microsoft.Network resource provider explicitly lists virtualNetworks as a movable resource.

CertEmpire



## Question: 60

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

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

You deploy a load balancer that has the following configurations:

Name: LB1

Type: Internal

SKU: Standard

Virtual network: VNET1

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

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

Does this meet the goal?

A. Yes

B. No

CertEmpire

**Answer:**

B

**Explanation:**

The proposed solution fails because of a SKU mismatch. A Standard SKU Azure Load Balancer requires that virtual machines in its backend pool either have no public IP address or have a Standard SKU public IP address. VM1 is configured with a Basic SKU public IP address, which makes it incompatible with the Standard SKU load balancer (LB1). The proposed solution of creating and associating another Basic SKU public IP address does not resolve this incompatibility. To meet the goal, VM1's Basic SKU public IP must be disassociated, or it must be upgraded to a Standard SKU.

**References:**

1. Azure Load Balancer SKUs Comparison: Microsoft Learn. "Virtual machines with a Standard SKU Public IP address or no Public IP address can be added to the backend pool of a Standard Load Balancer. Virtual machines with a Basic SKU Public IP address...can be added to the backend pool of a Basic Load Balancer." This explicitly states the compatibility rules.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/skus#skus>

2. Load Balancer and Public IP address SKUs: Microsoft Learn. "You can't have both basic and standard SKU resources. You can't mix SKU types for standalone virtual machines, availability sets, or virtual machine scale sets in the same backend pool." This reinforces the rule against mixing SKUs.

URL: <https://learn.microsoft.com/en-us/azure/load-balancer/skus#limitations>

CertEmpire