

DUMPSBOSS.

Microsoft Azure Administrator

Microsoft AZ-104

Version Demo

Total Demo Questions: 20

Total Premium Questions: 927

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

You need to meet the user requirement for Admin1.

What should you do?

- A. From the Azure Active Directory blade, modify the Groups
- B. From the Azure Active Directory blade, modify the Properties
- C. From the Subscriptions blade, select the subscription, and then modify the Access control (IAM) settings
- D. From the Subscriptions blade, select the subscription, and then modify the Properties

ANSWER: D

Explanation:

To meet the requirement for Admin1, you need to set them as the subscription's **Service administrator** (a classic admin role). That setting isn't managed through Azure AD group membership or normal RBAC roles in **Access control (IAM)**. IAM is for Azure RBAC roles like Owner/Contributor/Reader, but it doesn't change who the classic "Service admin" is.

The place where you actually change the Service administrator is at the subscription level. In the Azure portal, you go to **Subscriptions**, pick the right subscription, then open **Properties**. That's where you can view/change the Service admin so Admin1 can receive subscription-level notifications and communications tied to that classic role.

More details here: <https://learn.microsoft.com/en-us/azure/role-based-access-control/classic-administrators>

QUESTION NO: 2 - (DRAG DROP)

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

You need to back up VM. The solution must ensure that backups are stored across three availability zones in the primary region.

Which three actions 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.

The screenshot shows a drag-and-drop interface. On the left, under the heading "Actions", there is a list of five items, each in a light blue box with a right-pointing arrow on its right side:

- Set Replication to Zone-redundant storage (ZRS).
- Configure a replication policy.
- Set Replication to Locally-redundant storage (LRS).
- For VM1, create a backup policy and configure the backup.
- Create a Recovery Services vault.

On the right, under the heading "Answer Area", there are two empty slots, each with a left-pointing arrow on its left side, indicating where to drag the actions.

ANSWER:

Explanation:

To meet the requirement (backups stored across three availability zones in the *primary* region), you need Azure Backup to use **zone-redundant storage (ZRS)** for the Recovery Services vault. ZRS is the setting that keeps recovery points replicated across multiple availability zones within the same region (when the region and vault support it). If you don't set the vault to ZRS, the backups won't be guaranteed to be spread across zones.

Looking at the provided answer, the sequence chosen was: create a Recovery Services vault, configure VM backup policy, then configure a replication policy. That last step is the giveaway that the solution is off-track: a "replication policy" is something you'd configure for **Azure Site Recovery (ASR)** replication, not for Azure Backup of a VM. Azure VM Backup uses a **backup policy** (schedule + retention), and the storage redundancy is controlled at the vault level (LRS/GRS/ZRS depending on support). So the correct flow is: create the vault, set its storage redundancy to ZRS, then enable backup for VM1 using a backup policy.

Microsoft docs that back this up: ZRS for Recovery Services vault storage redundancy is described here: [Recovery Services vault overview](#) and the storage redundancy options (including ZRS where supported) are covered here: [Set storage redundancy for a Recovery Services vault](#). For enabling VM backup and applying a backup policy, see: [Back up an Azure VM in the Azure portal](#).

QUESTION NO: 3

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

Explanation:

A network interface (NIC) has to be created in the same Azure region as the virtual network (VNet) it will connect to. You can't create a NIC in one region and then attach it to a subnet in a VNet that lives in a different region.

From the resources in the table, the only region that has the required VNet available for NIC1 is East US. Even if you have other resources (like a VM or storage) in other regions, that doesn't help—what matters for the NIC is the VNet/subnet location.

So the only valid location you can pick for creating NIC1 is East US.

Reference: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-network-interface>

QUESTION NO: 4 - (HOTSPOT)

HOTSPOT

-

You have an Azure App Service web app named app1. You configure autoscaling as shown in following exhibit.

The screenshot shows the 'Auto created scale condition' configuration page for an Azure App Service. The page is titled 'Default*' and includes a 'Delete warning' section with an information icon and text: 'The very last or default recurrence rule cannot be deleted. Instead, you can disable autoscale to turn off autoscale.' Below this is the 'Scale mode' section with two radio buttons: 'Scale based on a metric' (selected) and 'Scale to a specific instance count'. The 'Rules' section contains an information icon and text: 'It is recommended to have at least one scale in rule. To create new rules, click Add a rule.' Underneath, there is a 'Scale out' section with a 'When' condition: '(Average) CpuPercentage > 70' and an 'Increase count by 1' value. A '+ Add a rule' button is visible. The 'Instance limits' section has three input fields: 'Minimum' with value '1', 'Maximum' with value '5', and 'Default' with value '1'. Each field has a dropdown arrow and a checkmark. At the bottom, there is a 'Schedule' section and a note: 'This scale condition is executed when none of the other scale condition(s) match'.

You configure the autoscale rule criteria as shown in the following exhibit.

Criteria

Time aggregation *

Maximum

Metric namespace * Metric name

1 minute time grain

Dimension Name	Operator	Dimension Values	Add
Instance	=	All values	+

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

CpuPercentage (Maximum)

1.67 %

Enable metric divide by instance count

Operator * Metric threshold to trigger scale action *

%

Duration (minutes) *

Time grain (minutes) Time grain statistic *

Time grain (minutes)

Time grain statistic *

Action

Operation *

Cool down (minutes) *

Instance count *

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

NOTE: Each correct selection is worth one point.

After CPU usage has reached 80 percent for 15 minutes, **[answer choice]** will be running.

- 1 instance
- 2 instances
- 3 instances
- 4 instances
- 5 instances

Once the first scale-out instance is created, the minimum time before an additional instance is created will be **[answer choice]**.

- 1 minute
- 5 minutes
- 10 minutes
- 15 minutes

ANSWER:

After CPU usage has reached 80 percent for 15 minutes, **[answer choice]** will be running.

- 1 instance
- 2 instances
- 3 instances
- 4 instances
- 5 instances

Once the first scale-out instance is created, the minimum time before an additional instance is created will be **[answer choice]**.

- 1 minute
- 5 minutes
- 10 minutes
- 15 minutes

Explanation:

Looking at the autoscale configuration in the exhibits, the App Service plan is set to scale based on a metric with instance limits of Minimum = 1, Default = 1, and Maximum = 5 (IMAGE_1). There's a single scale-out rule that increases the instance count by 1 when the CPU metric breaches the configured threshold.

The important part is the rule criteria (IMAGE_2). The rule is configured to trigger when **CPU Percentage** is **greater than 70%** for a **duration of 10 minutes**, and when it triggers it will **increase count by 1**. The rule also has a **cool down of 5 minutes**, which is the minimum wait time after a scale action before autoscale will perform another scale action. (In other words, even if the metric stays high, autoscale won't immediately add another instance until the cooldown has passed.)

So, if CPU usage reaches **80%** and stays there for **15 minutes**, that condition definitely satisfies the configured rule (because 80% is above the 70% threshold, and 15 minutes is longer than the 10-minute duration). Since the plan starts at 1 instance (Default = 1), the first scale-out action adds one instance, resulting in **2 instances** running.

After that first scale-out instance is created, the earliest autoscale could add another instance is after the **5-minute cooldown** period. This cooldown is specifically designed to prevent rapid "scale-out storms" and give the platform time to observe whether the previous scale action resolved the load.

References: [Understand autoscale settings](#), [Azure Monitor autoscale overview](#).

QUESTION NO: 5

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. NOTE: Each correct selection is worth one point.

- 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:

To add a site-to-site VPN as a backup path, you need the standard building blocks: a Local Network Gateway (to represent your on-premises VPN device and address spaces), a VPN Gateway in VNet1, and then a Connection that ties the two together.

You don't need to create a new gateway subnet here because VNet1 already connects via ExpressRoute. That setup typically requires a gateway subnet already, so creating another one would be redundant (and you can't have two gateway subnets anyway).

For cost, you'd normally think "Basic" SKU, but Basic VPN gateways can't be used in scenarios where the VNet also has ExpressRoute (you need a non-Basic SKU). So the cheapest suitable choice from the options is VpnGw1.

References: <https://learn.microsoft.com/en-us/azure/vpn-gateway/vpn-gateway-about-vpngateways> and <https://learn.microsoft.com/en-us/azure/expressroute/expressroute-howto-coexist-resource-manager>

QUESTION NO: 6

You plan to automate the deployment of a virtual machine scale set that uses the Windows Server 2016 Datacenter image.

You need to ensure that when the scale set virtual machines are provisioned, they have web server components installed.

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

NOTE: Each correct selection is worth one point.

- A. Upload a configuration script
- B. Create an automation account
- C. Create an Azure policy
- D. Modify the extensionProfile section of the Azure Resource Manager template
- E. Create a new virtual machine scale set in the Azure portal

ANSWER: A D

Explanation:

To install web server components automatically as VM Scale Set instances are provisioned, you should use a VM extension such as PowerShell DSC. This requires (1) having the DSC configuration available (for example, uploaded to a location the extension can access) and (2) configuring the scale set model (ARM template) to include the DSC extension in the `extensionProfile` so it runs during provisioning. References: <https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-dsc> , <https://learn.microsoft.com/en-us/azure/virtual-machines/extensions/dsc-windows>

QUESTION NO: 7 - (SIMULATION)

You need to configure Azure Backup to back up the file shares and virtual machines.

What is the minimum number of Recovery Services vaults and backup policies you should create? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

ANSWER: checkanswerinimageexplanation.

Explanation:

Explanation.

See the answer as below.



QUESTION NO: 8

You have an Azure Kubernetes Service (AKS) cluster named AKS1. You need to configure cluster autoscaler for AKS1.

Which two tools should you use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

A. the kubectl command

B. the az aks command

the az aks command: You can use the Azure Command-Line Interface (CLI) command `az aks update` to configure the cluster autoscaler for an AKS cluster. This command allows you to enable or disable the cluster autoscaler and set parameters like minimum and maximum node counts.

C. the Set-AzVm cmdlet

D. the Azure portal

the Azure portal: You can also configure the cluster autoscaler for AKS using the Azure portal. Navigate to your AKS cluster in the Azure portal, go to the "Node pools" section, and then configure the autoscaler settings for the specific node pool.

E. the Set-AzAks cmdlet

ANSWER: B D

Explanation:

To turn on (and tune) AKS cluster autoscaler, you typically do it at the AKS/node pool level, not from inside the cluster. That's why **Azure CLI** is a good fit: with `az aks nodepool update` (or `az aks update` in some cases), you can enable autoscaling and set the min/max node counts for a node pool.

You can also do the same thing straight from the **Azure portal**. In the AKS resource, go to *Node pools*, pick the pool, and enable autoscaling with the min/max values. It's basically the click-through version of the CLI settings.

`kubectl` is great for managing Kubernetes objects (deployments, services, etc.), but it doesn't configure the AKS-managed cluster autoscaler settings for node pools. And cmdlets like `Set-AzVm` are VM-focused, not AKS autoscaler configuration.

References: <https://learn.microsoft.com/en-us/azure/aks/cluster-autoscaler> and <https://learn.microsoft.com/en-us/azure/aks/scale-cluster>

QUESTION NO: 9

You have an Azure subscription that contains two virtual machines named VM1 and VM2

You create an Azure load balancer.

You plan to create a load balancing rule that will load balance HTTPS traffic between VM1 and VM2.

Which two additional load balance resources should you create before you can create the load balancing rule? Each correct answer presents part of the solution

MOTL Each correct selection 5 worth one point.

- A. a frontend IP address
- B. a backend pool
- C. a health probe
- D. an inbound NAT rule
- E. a virtual network

ANSWER: B C

Explanation:

Before you can create an HTTPS load-balancing rule, the load balancer needs to know two things: *where* to send traffic and *how* to tell if a VM is healthy.

The **backend pool** is the “where.” It’s the group that contains VM1 and VM2 (more specifically, their NICs/IP configs). Without a backend pool, there’s nowhere for the rule to forward incoming HTTPS requests.

The **health probe** is the “how.” Azure Load Balancer checks the probe to decide whether VM1 or VM2 should receive traffic. If a VM fails the probe, it gets taken out of rotation automatically, which prevents sending users to a broken endpoint.

A frontend IP is required for a rule, but it’s typically created when you set up the load balancer itself, not something you add afterward as an “additional” resource in this context. NAT rules are for direct inbound access (like RDP/SSH), not load balancing between two VMs.

References: <https://learn.microsoft.com/en-us/azure/load-balancer/load-balancer-components> and <https://learn.microsoft.com/en-us/azure/load-balancer/load-balancer-manage-rules>

QUESTION NO: 10

You have an Azure subscription that contains three virtual networks named VNet1, VNet2, VNet3. VNet2

contains a virtual appliance named VM2 that operates as a router.

You are configuring the virtual networks in a hub and spoke topology that uses VNet2 as the hub network.

You plan to configure peering between VNet1 and VNet2 and between VNet2 and VNet3.

You need to provide connectivity between VNet1 and VNet3 through VNet2.

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

NOTE: Each correct selection is worth one point.

- A. On the peering connections, allow forwarded traffic.
- B. On the peering connections, allow gateway transit.
- C. Create route tables and assign the table to subnets.
- D. Create a route filter.
- E. On the peering connections, use remote gateways.

ANSWER: A C

Explanation:

To enable spoke-to-spoke connectivity (VNet1 to VNet3) via a hub VNet (VNet2) that contains a router/NVA (VM2), you must: 1) Allow forwarded traffic on both peering links so that packets forwarded by the NVA in VNet2 are permitted to traverse the peering connections. 2) Use user-defined routes (route tables) associated to the spoke subnets to send traffic destined for the other spoke's address space to the NVA's IP in VNet2 (next hop = Virtual appliance). Without UDRs, Azure system routes do not provide transitive routing across VNet peerings. Official references: - VNet peering settings (including forwarded traffic): <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-network-peering-overview> - User-defined routes and virtual appliance next hop: <https://learn.microsoft.com/en-us/azure/virtual-network/virtual-networks-udr-overview>

QUESTION NO: 11

You need to deploy an Azure virtual machine scale set that contains five instances as quickly as possible. What should you do?

- A. Deploy five virtual machines. Modify the Availability Zones settings for each virtual machine.
- B. Deploy five virtual machines. Modify the Size setting for each virtual machine.
- C. Deploy one virtual machine scale set that is set to VM (virtual machines) orchestration mode.
- D. Deploy one virtual machine scale set that is set to ScaleSetVM orchestration mode.

Deploy one virtual machine scale set that is set to ScaleSetVM orchestration mode.

To deploy an Azure Virtual Machine Scale Set (VMSS) with five instances as quickly as possible, the best option is to use ScaleSetVM orchestration mode. This mode is specifically designed for scalable and automatically managed virtual machines.

ANSWER: D

Explanation:

To get five instances out the door fast, you should deploy a single VM scale set in **ScaleSetVM** orchestration mode. That mode is the “classic” VMSS experience where Azure treats the instances as a uniform set and can create them in parallel using the scale set model.

By contrast, **VM** orchestration mode (flexible orchestration) is meant for more advanced scenarios where instances behave more like individually managed VMs (mixing sizes, tighter per-VM control, etc.). It’s great when you need that flexibility, but it’s not the simplest/quickest path when your only goal is “give me five identical instances ASAP.”

Also, deploying five separate VMs and then tweaking settings is slower and more manual than letting VMSS stamp out the instances for you in one deployment.

References: <https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-orchestration-modes> and <https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/overview>

QUESTION NO: 12

You have an Azure subscription named Subscription1 that has the following providers registered:

- Authorization
- Automation
- Resources
- Compute
- KeyVault
- Network
- Storage
- Billing ▪ Web

Subscription1 contains an Azure virtual machine named VM1 that has the following configurations:

- Private IP address: 10.0.0.4 (dynamic)
- Network security group (NSG): NSG1
- Public IP address: None
- Availability set: AVSet
- Subnet: 10.0.0.0/24
- Managed disks: No ▪ Location: East US

You need to record all the successful and failed connection attempts to VM1.

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

NOTE: Each correct selection is worth one point.

- A.** Enable Azure Network Watcher in the East US Azure region.
- B.** Add an Azure Network Watcher connection monitor.

- C. Register the MicrosoftLogAnalytics provider.
- D. Create an Azure Storage account.
- E. Register the Microsoft.Insights resource provider.
- F. Enable Azure Network Watcher flow logs.

ANSWER: A E F

Explanation:

To record both allowed (successful) and denied (failed) connection attempts, the right tool is **NSG flow logs**. They capture traffic decisions made by the NSG (NSG1 in this case), so you can see what was permitted and what got blocked.

First, NSG flow logs are a feature of **Azure Network Watcher**, and Network Watcher is enabled per region. Since VM1 is in **East US**, you need Network Watcher turned on there, otherwise you won't be able to enable flow logs for NSG1.

Next, flow logs rely on the **Microsoft.Insights** resource provider (it's used by Network Watcher logging/metrics features). If it's not registered, enabling flow logs can fail or the feature won't be available as expected.

Finally, you enable the actual **flow logs** setting. (In real life you also pick a storage destination, but the question only asks for three actions.) For step-by-step details, see <https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-nsg-flow-logging-portal> and the overview at <https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-overview>.

QUESTION NO: 13 - (DRAG DROP)

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

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

Remove peering between VNet1 and VNet2.

Add the 10.33.0.0/16 address space to VNet1.

Recreate peering between VNet1 and VNet2.



Explanation:

The correct sequence is exactly what's shown in the answer area: first remove the peering between VNet1 and VNet2, then add the new 10.33.0.0/16 address space to VNet1, and finally recreate the peering. The key detail is that Azure blocks address space changes on a virtual network while it's peered. So even though adding an address space sounds like a simple configuration update, Azure requires you to break the peering relationship first. Once the peering is removed, you can safely add 10.33.0.0/16 to VNet1's address space. After that, you recreate the peering so that VNet1 and VNet2 can communicate again.

It's also important to notice what you *don't* need here. You don't need to delete VNet1 or create a new VNet, because the requirement is to extend VNet1's address space, not replace the network. You also don't need to enable "allow gateway transit" on either peering unless you're specifically trying to share a VPN/ExpressRoute gateway across peerings. The question only asks to ensure hosts on VNet1 and VNet2 can communicate after the address space change, which standard peering already provides once recreated.

Microsoft documents this limitation and the required workflow (delete peering → change address space → recreate peering) in the VNet peering management guidance: [Manage virtual network peering](#). You can also review general peering behavior and requirements here: [Virtual network peering overview](#).

QUESTION NO: 14

You have an Azure subscription named Subscription1. Subscription1 contains the resource groups in the following table.

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

RG1 has a web app named WebApp1. WebApp1 is located in West Europe.

You move WebApp1 to RG2.

What is the effect of the move?

- A. The App Service plan for WebApp1 remains in West Europe. Policy2 applies to WebApp1.
- B. The App Service plan for WebApp1 moves to North Europe. Policy2 applies to WebApp1.
- C. The App Service plan for WebApp1 remains in West Europe. Policy1 applies to WebApp1.
- D. The App Service plan for WebApp1 moves to North Europe. Policy1 applies to WebApp1.

ANSWER: A

Explanation:

Moving an Azure Web App to a different resource group doesn't "relocate" it to a new region. App Service apps run in the region of their App Service plan, and a move between resource groups is basically just changing the management container (the RG), not the hosting location. So WebApp1 will still be running on the same App Service plan in West Europe after the move.

The other key change is governance: Azure Policy assignments applied at the resource-group scope follow the resource group. Once WebApp1 is in RG2, the policies assigned to RG2 are the ones that will evaluate and apply to WebApp1. That's why Policy2 (the one on RG2) becomes the effective policy after the move.

References: <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/move-resource-group-and-subscription> and <https://learn.microsoft.com/en-us/azure/app-service/app-service-plan-manage>.

QUESTION NO: 15

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

[Create new](#)

INSTANCE DETAILS

* Virtual machine name ⓘ VM1

* Region ⓘ (US) West US 2

Availability options ⓘ No infrastructure redundancy required

* Image ⓘ Windows Server 2016 Datacenter

[Browse all public and private images](#)

Azure Spot instance ⓘ Yes No

* Size ⓘ **Standard DS1 v2**
1 vcpu, 3.5 GiB memory (ZAR 632.47/month)
[Change size](#)

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

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.

i 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
[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. NOTE: Each correct selection is worth one point.

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

ANSWER: A C

Explanation:

To place a VM into an Availability Zone, you have to pick a zonal deployment option during VM creation. That's done under the VM's **Availability options** (sometimes shown as "Availability zone" in the portal). If you leave this as "No infrastructure redundancy required" or pick an availability set instead, the VM won't be created in a specific zone.

You also need to use **managed disks**. Zonal VMs rely on managed disks because Azure needs to pin the disks to the same zone as the VM. If you were using unmanaged disks (storage accounts), you can't guarantee zone placement, and Azure won't let you create a zonal VM with that setup.

Settings like VM size, image, or OS disk type can matter for other reasons (for example, whether a size is available in a zone), but they aren't the core "must change" settings implied here. The two direct knobs you change to make the VM zonal are choosing the Availability Zone option and ensuring disks are managed.

References: <https://learn.microsoft.com/en-us/azure/virtual-machines/availability-zone-overview> and <https://learn.microsoft.com/en-us/azure/virtual-machines/disks-types>

QUESTION NO: 16

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

Name	Azure region
VM1	West Europe
VM2	West Europe
VM3	North Europe
VM4	North Europe

You have a Recovery Services vault that protects VM1 and VM2.

You need to protect VM3 and VM4 by using Recovery Services.

What should you do first?

- A. Create a new Recovery Services vault
- B. Create a storage account
- C. Configure the extensions for VM3 and VM4
- D. Create a new backup policy

ANSWER: A

Explanation:

The first thing to check with Azure Backup is location: a Recovery Services vault can only back up Azure VMs that are in the same Azure region as the vault. Since the existing vault already protects VM1 and VM2, it's almost certainly sitting in their region. If VM3 and VM4 are in a different region (as shown in the table), you can't just add them to the same vault.

So the right first step is to create a new Recovery Services vault in the region where VM3 and VM4 live. After that, you can enable backup for each VM and pick (or create) a backup policy as needed. You don't need to create a storage account for this, and you typically don't manually install extensions first—the backup process handles the required VM extension when you enable protection.

Microsoft calls this out in the VM backup support matrix and the VM backup setup docs: <https://learn.microsoft.com/en-us/azure/backup/backup-support-matrix-iaas#support-for-azure-vm-backup> and <https://learn.microsoft.com/en-us/azure/backup/backup-azure-vms-first-look-arm>.

QUESTION NO: 17 - (HOTSPOT)

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

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

You plan to create a data collection rule named DCRI in Azure Monitor.

Which resources can you set as data sources in DCRI, and which resources can you set as destinations in DCRI? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Data sources:

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

Destinations:

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

ANSWER:

Answer Area

Data sources:

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

Destinations:

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

Explanation:

Looking at the resources you have (VM1, storage1, Workspace1, and DB1), the key is to remember what a data collection rule (DCR) actually does in Azure Monitor. A DCR defines *what* data to collect (like Windows event logs, syslog, and performance counters) and *where* to send it. In most AZ-104 scenarios, that collection happens through the Azure Monitor Agent (AMA), which runs on a machine (for example, an Azure VM). That's why VM1 is the only resource in the table that makes sense as a DCR data source: it's the only one that can run the agent and generate those OS-level telemetry streams.

On the destination side, DCRs commonly send collected log data to a Log Analytics workspace. In your list, that's Workspace1. A storage account (storage1) is a common place to store many kinds of data in Azure, but it isn't the expected/standard DCR destination for AMA-based log collection in this exam-style question. Likewise, an Azure SQL database (DB1) isn't used as a DCR destination for Azure Monitor agent-collected logs. So the correct pairing is: data sources = VM1 only, destinations = Workspace1 only.

You can confirm the role of DCRs and their relationship to AMA and Log Analytics in Microsoft's documentation for data collection rules and Azure Monitor Agent: [Data collection rules overview](#) and [Azure Monitor Agent overview](#).

QUESTION NO: 18

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

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

You have an Azure Directory (Azure AD) tenant named Adatum and an Azure Subscription named Subscription1. Adatum contains a group named Developers. Subscription1 contains a resource group named Dev.

You need to provide the Developers group with the ability to create Azure logic apps in the Dev resource group.

Solution: On Dev, you assign the Logic App Contributor role to the Developers group.

Does this meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes, this meets the goal. If you assign the **Logic App Contributor** role to the Developers group at the **Dev resource group** scope, they get the rights needed to create and manage Logic Apps inside that resource group.

The nice part is that it's scoped correctly: they can work with Logic Apps in Dev without being given broad permissions like full Contributor on every resource type, or Owner permissions. That keeps access tighter while still letting them do their job.

You can double-check what the built-in role allows in Microsoft's RBAC built-in roles list here: <https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>.

QUESTION NO: 19

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

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

You deploy an Azure Kubernetes Service (AKS) cluster named AKS1. You need to deploy a YAML file to AKS1.

Solution: From Azure CLI, you run azcopy. Does this meet the goal?

A. Yes

B. No

ANSWER: B

Explanation:

No. `azcopy` is a file transfer tool (typically used for copying data to and from Azure Storage), not a Kubernetes deployment tool. Running `azcopy` won't apply Kubernetes manifests to an AKS cluster, so it doesn't deploy the YAML in the way AKS expects.

To deploy a YAML manifest to AKS, you normally connect to the cluster and use `kubectl`, for example `kubectl apply -f yourfile.yaml`. That command tells the Kubernetes API server to create or update the resources defined in the YAML (like Deployments, Services, Ingress, and so on).

References: <https://learn.microsoft.com/en-us/azure/aks/kubernetes-walkthrough> and <https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10>

QUESTION NO: 20 - (SIMULATION)

You have an Azure App Service web app named `app1`.

You configure autoscaling as shown in following exhibit.

The screenshot shows the 'Auto created scale condition' configuration page for an Azure App Service web app. The page is titled 'Default*' and includes several sections:

- Delete warning:** A message stating that the very last or default recurrence rule cannot be deleted, but autoscale can be disabled.
- Scale mode:** Two radio buttons are present: 'Scale based on a metric' (which is selected) and 'Scale to a specific instance count'.
- Rules:** A message recommends having at least one scale rule. Below this, a 'Scale out' rule is defined with the condition '(Average) CpuPercentage > 70' and the action 'Increase count by 1'. There is an '+ Add a rule' button.
- Instance limits:** Three input fields are shown: 'Minimum' (set to 1), 'Maximum' (set to 5), and 'Default' (set to 1). Each field has a dropdown arrow and a checkmark.
- Schedule:** A note at the bottom states 'This scale condition is executed when none of the other scale condition(s) match'.

You configure the autoscale rule criteria as shown in the following exhibit.

The screenshot shows the configuration for a scaling policy. The 'Criteria' section includes a line graph of CPU usage over time, with a threshold line at 80%. Below the graph, the 'CpuPercentage (Maximum)' is set to 1.67%. The 'Operator' is 'Greater than' and the 'Metric threshold to trigger scale action' is 70%. The 'Duration (minutes)' is 10, and the 'Time grain (minutes)' is 1. The 'Time grain statistic' is 'Average'. The 'Action' section shows the 'Operation' is 'Increase count by' and the 'Instance count' is 1.

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

Answer Area

After CPU usage has reached 80 percent for 15 minutes, [answer choice] will be running.

Once the first scale-out instance is created, the minimum time before an additional instance is created will be [answer choice].

ANSWER: Seetheanswerbelow.

Explanation:

Answer is below in image.

Answer Area

After CPU usage has reached 80 percent for 15 minutes, [answer choice] will be running.

Once the first scale-out instance is created, the minimum time before an additional instance is created will be [answer choice].