



- Expert Verified, Online, **Free**.



## **CERTIFICATION TEST**

- [CertificationTest.net](https://CertificationTest.net) - Cheap & Quality Resources With Best Support

You have two Hyper-V hosts named Host1 and Host2. Host1 has an Azure virtual machine named VM1 that was deployed by using a custom Azure Resource Manager template. You need to move VM1 to Host2. What should you do?

- A. From the Update management blade, click Enable.
- B. From the Overview blade, move VM1 to a different subscription.
- C. From the Redeploy blade, click Redeploy.
- D. From the Profile blade, modify the usage location.

**Suggested Answer: C**

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.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/redeploy-to-new-node>

Community vote distribution

C (100%)

🗳️ **brenobaronte** 1 month, 1 week ago

**Selected Answer: C**

C is correct.

upvoted 3 times

🗳️ **Mirdwal** 1 month, 4 weeks ago

**Selected Answer: C**

C. From the Redeploy blade, click Redeploy.

Why the Other Options Are Wrong:

A. From the Update management blade, click Enable

This is for patching VMs, not moving them.

B. From the Overview blade, move VM1 to a different subscription

That moves the VM to another subscription, not to another host.

D. From the Profile blade, modify the usage location

Usage location is for licensing and billing (not VM placement).

upvoted 4 times

🗳️ **anversadat** 2 months, 3 weeks ago

**Selected Answer: B**

my option

upvoted 1 times

🗳️ **Azurite** 3 months, 1 week ago

**Selected Answer: C**

Answer C: (Closest)

But here is my understanding.

None of the provided options seems to align directly with the physical migration of a VM from one Hyper-V host (Host1) to another (Host2). Azure's native management options are tailored to managing VMs within Azure's infrastructure rather than facilitating migration to on-premises Hyper-V hosts.

Thus, the given answers might be testing our ability to discern which action is most contextually appropriate within Azure's environment, even if none directly achieves moving the VM to Host2. Among these, option C ("Redeploy") seems closest, as it involves "moving" the VM within Azure, but it does

not resolve the cross-environment migration requirement.

Is my understanding correct?

upvoted 1 times

🗳️ 👤 **starasenko** 3 months, 4 weeks ago

**Selected Answer: C**

C - is a correct answer, as the Redeploy operation will change the Host. If you want to restart the VM without changing the Host, just use Restart operation

upvoted 1 times

🗳️ 👤 **Aidan\_Toh** 9 months, 1 week ago

Answer C is correct.

upvoted 1 times

🗳️ 👤 **forgetfulalligator** 9 months, 1 week ago

**Selected Answer: C**

The "Redeploy" blade in the Azure portal can be used to redeploy a virtual machine to a different host. When you redeploy a virtual machine, it moves to a new host within the same Azure datacenter

upvoted 3 times

🗳️ 👤 **james2033** 9 months, 1 week ago

**Selected Answer: C**

Redeploy blade <https://learn.microsoft.com/en-us/troubleshoot/azure/virtual-machines/redeploy-to-new-node-windows#use-the-azure-portal>

upvoted 2 times

🗳️ 👤 **Rhashik** 11 months, 2 weeks ago

**Selected Answer: C**

Answer is C

upvoted 1 times

🗳️ 👤 **scotttohoty** 1 year, 1 month ago

**Selected Answer: C**

agree C

upvoted 1 times

🗳️ 👤 **collocollo** 1 year, 9 months ago

**Selected Answer: C**

C - ok

upvoted 1 times

🗳️ 👤 **ndoan** 1 year, 10 months ago

agree! C is the right answer

upvoted 2 times

🗳️ 👤 **ITRogers** 1 year, 11 months ago

C is the right answer

upvoted 1 times

🗳️ 👤 **torsanta** 2 years ago

**Selected Answer: C**

Chosen Answer: C

upvoted 1 times

🗳️ 👤 **Kallem\_Niharika** 2 years, 1 month ago

C is the correct answer

upvoted 1 times

🗳️ 👤 **nethala\_swathi** 2 years, 1 month ago

Option:C

upvoted 1 times

🗳️ 👤 **D3D1997** 2 years, 1 month ago

**Selected Answer: C**

redeploy

upvoted 2 times

DRAG DROP -

You have downloaded an Azure Resource Manager template to deploy numerous virtual machines. The template is based on a current virtual machine, but must be adapted to reference an administrative password.

You need to make sure that the password is not stored in plain text.

You are preparing to create the necessary components to achieve your goal.

Which of the following should you create to achieve your goal? Answer by dragging the correct option from the list to the answer area.

Select and Place:

## Options

## Answer

An Azure Key Vault

An Azure Storage account

Azure Active Directory (AD)  
Identity Protection

An access policy

An Azure policy

A backup policy



## Options

## Answer

An Azure Storage account


Azure Active Directory (AD)  
Identity Protection

An Azure policy

A backup policy

An Azure Key Vault

An access policy


 **vincyN90** Highly Voted 2 years, 9 months ago  
Key Vault + Access Policy.


Using Key Vault we create a secret containing our Password: <https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-portal> .


Using an Access Policy we allow access to the previously created secret.

Documentation Guide: <https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/key-vault-parameter?tabs=azure-cli>  
upvoted 43 times

 **moned** Highly Voted 2 years, 2 months ago  
Is it me, or is the question misleading since it says "option" - singular?  
upvoted 10 times

 **Rene\_Clevers** 2 years, 1 month ago  
I had the same problem  
upvoted 4 times

 **brenobaronte** Most Recent 1 month, 1 week ago  
Key Vault + Access Policy.  
upvoted 1 times

 **Azurite** 3 months, 1 week ago  
Key Vault + Access Policy

Azure Key Vault is the primary service for securely storing and managing sensitive information like passwords, keys, and secrets. However, simply storing the password in Key Vault isn't sufficient—permissions also need to be configured so that only authorized resources or users can access the secret. This is where the access policy comes into play.

Azure Key Vault is the primary service for securely storing and managing sensitive information like passwords, keys, and secrets. However, simply storing the password in Key Vault isn't sufficient—permissions also need to be configured so that only authorized resources or users can access the secret. This is where the access policy comes into play.

upvoted 1 times

🗨️ 👤 **samtestking** 5 months, 2 weeks ago

Test comment

upvoted 2 times

🗨️ 👤 **man5484** 9 months, 1 week ago

To ensure that the password is not stored in plain text in an Azure Resource Manager template, you should utilize an Azure Key Vault. Azure Key Vault provides a secure way to store and manage sensitive information such as passwords, cryptographic keys, and certificates. By storing the password in Azure Key Vault, it remains encrypted and can be securely accessed by authorized services or applications.

upvoted 1 times

🗨️ 👤 **nivramarvino** 9 months, 1 week ago

The user just needs to reference the administrative password, right? Am I understanding this correctly?

"The access policies aren't needed if the user is deploying a template that retrieves a secret. Add a user to the access policies only if the user needs to work

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/key-vault-parameter?tabs=azure-cli#:~:text=The%20access%20policies%20aren%27t%20needed%20if%20the%20user%20is%20deploying%20a%20template%20that%20retrieves%20a%20secret>

upvoted 1 times

🗨️ 👤 **cemeti** 10 months, 3 weeks ago

Using an Access Policy we allow access to the previously created secret.

Documentation Guide: <https://sites.google.com/view/learnmicrosoftcomenustrainingm/home>

upvoted 1 times

🗨️ 👤 **mpxl** 1 year, 1 month ago

They might update it to Key Vault + RBAC, because access policy is deprecated now

upvoted 5 times

🗨️ 👤 **dimipap3** 1 year, 8 months ago

Key Vault + Access Policy, Is the correct but the Access Policy option is legacy now, and most likely will be replaced with the option of Azure role-based access control instead of Access Policy

upvoted 5 times

🗨️ 👤 **SmartTechieGrv** 2 years, 2 months ago

Key Vault and Access Policy both

upvoted 2 times

🗨️ 👤 **p2006** 2 years, 4 months ago

<https://learn.microsoft.com/en-us/azure/key-vault/general/assign-access-policy>

upvoted 1 times

🗨️ 👤 **forgetfulalligator** 2 years, 4 months ago

Azure Key Vault is the recommended approach for securely storing and managing secrets such as the administrative password used in an Azure Resource Manager template. However, an access policy can also be used to control access to a secret stored in a different Azure resource, such as an Azure Storage account.

upvoted 1 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

Key Vault & Access Policy

upvoted 1 times

🗨️ 👤 **Santmax** 2 years, 5 months ago

Key Vault + Access Policy.

upvoted 1 times

🗨️ 👤 **Obioha** 3 years, 2 months ago

<https://www.examttopics.com/exams/microsoft/az-104/view/2/>

Answer is correct, see question in AZ104 #20

upvoted 5 times

Your company has an Azure Kubernetes Service (AKS) cluster that you manage from an Azure AD-joined device. The cluster is located in a resource group.

Developers have created an application named MyApp. MyApp was packaged into a container image.

You need to deploy the YAML manifest file for the application.

Solution: You install the Azure CLI on the device and run the `kubectl apply -f myapp.yaml` command.

Does this meet the goal?

A. Yes

B. No

**Suggested Answer: A**

`kubectl apply -f myapp.yaml` applies a configuration change to a resource from a file or stdin.

Reference:



<https://kubernetes.io/docs/reference/kubectl/overview/>

<https://docs.microsoft.com/en-us/cli/azure/aks>

Community vote distribution

A (85%)

B (15%)

  **guschess** Highly Voted 2 years, 5 months ago

# Connect to the cluster

To manage a Kubernetes cluster, use the Kubernetes command-line client, `kubectl`. `kubectl` is already installed if you use Azure Cloud Shell.

- Install `kubectl` locally using the `az aks install-cli` command:

Azure CLI

- `az aks install-cli`

Why answer A is correct?

url: <https://learn.microsoft.com/en-us/azure/aks/learn/quick-kubernetes-deploy-cli>

thanks

upvoted 8 times

  **PaoLoi** 2 months ago

Yes this is correct

upvoted 1 times

  **kingalexp** Most Recent 1 month ago

**Selected Answer: B**

B. Syntax error (`kubectl apply -f myapp.yaml`)

upvoted 1 times

  **brenobaronte** 1 month, 1 week ago

**Selected Answer: A**

A is the correct one.

upvoted 1 times

  **adelaTM** 3 months ago

**Selected Answer: B**

The answer is `No`, because there is a typo in the command; it should be with `-f`:

`kubectl apply -f myapp.yaml`

upvoted 3 times

  **Ira89** 3 months, 2 weeks ago

**Selected Answer: A**

kubectl apply -f myapp.yaml This is correct command to run manifest

upvoted 1 times

  **HernanP** 4 months, 1 week ago

**Selected Answer: B**

B

# Create a service using the definition in example-service.yaml.

kubectl apply -f example-service.yaml

# Create a replication controller using the definition in example-controller.yaml.

kubectl apply -f example-controller.yaml

# Create the objects that are defined in any .yaml, .yml, or .json file within the <directory> directory.

kubectl apply -f <directory>


upvoted 1 times

  **DheerajPatyal** 5 months ago

**Selected Answer: B**

No, this solution does not meet the goal as stated because the command provided contains a syntax error. The correct kubectl apply command should use the -f flag instead of "f

upvoted 1 times

  **ErickIMT** 5 months, 1 week ago

**Selected Answer: B**

The response must be "B", because the command is not correct at all, it says "kubectl apply `f myapp.yaml command" but must be "kubectl apply -f myapp.yaml command"

upvoted 2 times

  **Cicy** 5 months, 2 weeks ago

**Selected Answer: A**

Installing the Azure CLI enables you to interact with Azure resources, including AKS.

upvoted 1 times

  **Nhantran20000502** 8 months, 1 week ago

**Selected Answer: A**

By installing the Azure CLI and running the kubectl apply -f myapp.yaml command, you are correctly deploying the YAML manifest file for the application to the AKS cluster<sup>12</sup>. This command parses the manifest file and creates the defined Kubernetes objects in your AKS cluster.

upvoted 2 times

  **CESBCN** 8 months, 3 weeks ago

Option - B


"Solution: You install the Azure CLI on the device and run the kubectl apply `f myapp.yaml command."

Before run the kubectl on the device, you need to install kubectl.

"To manage a Kubernetes cluster, use the Kubernetes command-line client, kubectl. kubectl is already installed if you use Azure Cloud Shell. To install kubectl locally, use the az aks install-cli command. "

<https://learn.microsoft.com/en-us/azure/aks/learn/quick-kubernetes-deploy-cli#connect-to-the-cluster>

upvoted 3 times

  **man5484** 9 months, 1 week ago

**Selected Answer: A**

A

The solution provided involves using kubectl, which is the Kubernetes command-line tool, to apply the YAML manifest file (myapp.yaml) for deploying the application. Since you're managing the AKS cluster from an Azure AD-joined device and have the Azure CLI installed, running kubectl apply -f myapp.yaml command will deploy the application to the AKS cluster. This is a common and valid method for deploying applications to Kubernetes clusters, including AKS. Therefore, option A is correct.

upvoted 2 times

  **Amuktha** 1 year, 3 months ago

option B.

Azure CLI might help you to connect AKS but to run kubectl commands you need to install both kubectl and kubelogin

upvoted 1 times

  **mrsmjparkar** 1 year ago

However, you can also use Azure Cloud Shell and would not then need to install anything as it is built into cloud shell. I do it all the time.

upvoted 1 times

🗋️ 👤 **Mohanraj\_SV5** 1 year, 6 months ago

**Selected Answer: A**

A - Correct

upvoted 1 times

🗋️ 👤 **RikinPatel** 1 year, 8 months ago

Ans.: B

As per my understanding. step as follow as per microsoft..

- Create a resource group
- Create an AKS cluster
- Connect to the cluster
- Deploy the application

<https://learn.microsoft.com/en-us/azure/aks/learn/quick-kubernetes-deploy-cli>

Missing step: Connect to the cluster

upvoted 1 times

🗋️ 👤 **MikeAWS** 1 year, 9 months ago

A should be correct!

kubectl apply -f myapp.yaml

upvoted 1 times

🗋️ 👤 **collocollo** 1 year, 9 months ago

**Selected Answer: A**

A - ok

upvoted 1 times

Your company has an Azure Kubernetes Service (AKS) cluster that you manage from an Azure AD-joined device. The cluster is located in a resource group.

Developers have created an application named MyApp. MyApp was packaged into a container image.

You need to deploy the YAML manifest file for the application.

Solution: You install the docker client on the device and run the `docker run -it microsoft/azure-cli:0.10.17` command.

Does this meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Community vote distribution


B (100%)

 **jay158** Highly Voted 3 years, 12 months ago

Correct Answer NO

`docker run -it microsoft/azure-cli:0.10.17` is no k8s command

upvoted 19 times

 **lugospod** Highly Voted 3 years, 5 months ago

K8 Out of scope for 2021!

upvoted 12 times

 **Ira89** Most Recent 3 months, 2 weeks ago

**Selected Answer: B**

Kubernetes don't use docker command.


upvoted 1 times

 **Nhantran20000502** 8 months, 1 week ago

**Selected Answer: B**

Installing the Docker client and running the `docker run -it microsoft/azure-cli:0.10.17` command will start an interactive Azure CLI session within a Docker container, but it does not directly deploy the YAML manifest file to the AKS cluster. To deploy the YAML manifest file, you need to use the `kubectl apply -f myapp.yaml` command, which is specifically designed for interacting with Kubernetes clusters

upvoted 3 times

 **L\_\_L** 1 year, 2 months ago

**Selected Answer: B**

Kubernetes cannot use docker command.

upvoted 1 times


 **man5484** 1 year, 2 months ago

**Selected Answer: B**

The provided solution is incorrect for deploying the YAML manifest file for the application (MyApp) to the Azure Kubernetes Service (AKS) cluster. Using the Docker client (`docker run -it microsoft/azure-cli:0.10.17` command) would only allow you to interact with Docker containers, not Kubernetes clusters.

To deploy the YAML manifest file for the application to the AKS cluster, you should use the Kubernetes command-line tool (`kubectl`) rather than the Docker client. Therefore, the solution provided does not meet the goal.

upvoted 1 times

 **Ciupaz** 1 year, 5 months ago

Is a Docker command, not a Kubernetes command. So answer is No.

upvoted 1 times

 **Mohanraj\_SV5** 1 year, 6 months ago

**Selected Answer: B**

B - No

upvoted 1 times

🗄️ 👤 **colloccollo** 1 year, 9 months ago

**Selected Answer: B**

B - No

upvoted 1 times

🗄️ 👤 **viking2** 2 years ago

**Selected Answer: B**

NO, use kubectl

upvoted 1 times

🗄️ 👤 **SmartTechieGrv** 2 years, 2 months ago

Correct Answer No

upvoted 1 times

🗄️ 👤 **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: B**

To apply the Kubernetes manifest file to the AKS cluster, you would typically use the kubectl apply command. This command applies the desired state described in the manifest file to the cluster, which will result in the creation and deployment of the specified resources, including the containerized application.

upvoted 3 times

🗄️ 👤 **uffuchsi** 2 years, 4 months ago

B - No

upvoted 1 times

🗄️ 👤 **guschess** 2 years, 5 months ago

**Selected Answer: B**

The answer 3 helps with this topic.

upvoted 1 times

🗄️ 👤 **ExamTopicsAnkur** 2 years, 6 months ago

**Selected Answer: B**

Answer is No

upvoted 2 times

🗄️ 👤 **Chouffie** 2 years, 8 months ago

The Answer is NO

To deploy a YAML file to Kubernetes you need to use the "kubectl" command. An example could be: kubectl apply -f myapp.yaml

upvoted 5 times

🗄️ 👤 **HafizSalmanMalik** 2 years, 8 months ago

Correct answer is No because " docker run -it microsoft/azure-cli:0.10.17 " it's a Docker command, not a K8's command. Thanks

upvoted 2 times



Your company has a web app named WebApp1.

You use the WebJobs SDK to design a triggered App Service background task that automatically invokes a function in the code every time new data is received in a queue.

You are preparing to configure the service processes a queue data item.

Which of the following is the service you should use?


- A. Logic Apps
- B. WebJobs
- C. Flow
- D. Functions

**Suggested Answer: B**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-compare-logic-apps-ms-flow-webjobs>

Community vote distribution



**DefaultName2** Highly Voted 9 months, 1 week ago

"You use the WebJobs SDK to design a triggered App Service background task"

B - WebJob

Usually you'll host the WebJobs SDK in Azure WebJobs, but you can also run your jobs in a Worker Role. The Azure WebJobs feature of Azure Web Apps provides an easy way for you to run programs such as services or background tasks in a Web App...

<https://github.com/Azure/azure-webjobs-sdk>

upvoted 26 times

**eazybrave** 3 years, 2 months ago

Good sir/ma can you please share the material I can use to prepare for this exam. I will appreciate it thanks

upvoted 3 times

**Evo\_Morales** 3 years, 1 month ago

It is not specifically spelled out anywhere. There are only references to the types of things you need to know and suggested resources to visit.

upvoted 1 times

**coffecold** 2 years, 8 months ago

Azure WebJobs is not a service, you'll host the WebJobs SDK functions in Azure WebApps.

With Webjobs you program code in your webapp and you'll be able to execute 'kind of Azure functions' inside your Web App. It is automatically deployed with the deployment of the web app. WebJobs is the predecessor of Azure functions.

upvoted 7 times

**tracy161** 2 years, 5 months ago

Azure Webjobs with the Webjobs SDK is a code-first integration service that is built on Azure App Service and support features such as source control integration, authentication, and monitoring with Application Insights integration.

upvoted 1 times

**coffecold** 2 years, 8 months ago

But still the best answer is B.

upvoted 1 times

**noip** Highly Voted 3 years, 11 months ago

the Correct Answer, D: Function

upvoted 10 times

**edengoforit** 3 years, 5 months ago

Without any rationale?

upvoted 7 times

**Mega\_GA** Most Recent 1 month, 1 week ago

Selected Answer: D

D Functions

Azure Functions is built on the WebJobs SDK, so it shares many of the same event triggers and connections to other Azure services.

See <https://learn.microsoft.com/en-us/azure/azure-functions/functions-compare-logic-apps-ms-flow-webjobs#comparison-table>

upvoted 1 times

🗨️ 👤 **Mega\_GA** 1 month, 1 week ago

Selected Answer: C

Azure Functions is built on the WebJobs SDK, so it shares many of the same event triggers and connections to other Azure services. Some factors to consider when you're choosing between Azure Functions and WebJobs with the WebJobs SDK (See: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-compare-logic-apps-ms-flow-webjobs#comparison-table>)

upvoted 1 times

🗨️ 👤 **Thulas** 1 month, 2 weeks ago

Selected Answer: D

The correct answer is: D. Functions

Explanation:

The scenario describes the use of the WebJobs SDK to automatically invoke a function in the code every time new data is received in a queue. This behavior aligns exactly with Azure Functions, especially when combined with the WebJobs SDK. Azure Functions is built on top of the WebJobs SDK and provides a serverless compute platform designed for such event-driven tasks.

upvoted 1 times

🗨️ 👤 **ankita0910** 4 months, 1 week ago

Selected Answer: D

While WebJobs is a powerful feature for running background tasks within Azure App Service, Azure Functions is generally preferred for processing queue data items due to several advantages:

upvoted 1 times

🗨️ 👤 **ankita0910** 4 months, 1 week ago

Selected Answer: D

Azure Functions is built on the WebJobs SDK.

Azure Functions version 2.x is built on WebJobs SDK version 3.x.

Azure Functions version 1.x is built on WebJobs SDK version 2.x.

upvoted 1 times

🗨️ 👤 **Ajju5791** 6 months ago

Selected Answer: B

Azure Webjobs: scheduled, Azure function: respond to the demand

upvoted 1 times

🗨️ 👤 **Nhantran20000502** 8 months, 1 week ago

Selected Answer: B

The WebJobs SDK is specifically designed to run background tasks within an Azure Web App. It allows you to create triggered functions that respond to events such as new data in a queue. This makes it the appropriate service for processing queue data items in your scenario.

upvoted 1 times

🗨️ 👤 **CESBCN** 8 months, 3 weeks ago

Selected Answer: B

"WebJobs is a feature of Azure App Service that enables you to run a program or script in the same instance as a web app. All app service plans support WebJobs. There's no extra cost to use WebJobs."

<https://learn.microsoft.com/en-us/azure/app-service/webjobs-create?tabs=windowscode>

upvoted 1 times

🗨️ 👤 **Jobalos009** 9 months, 1 week ago

Selected Answer: B

Correct answer: B,

<https://learn.microsoft.com/en-us/azure/app-service/webjobs-create?tabs=windowscode> (Run background tasks with WebJobs in Azure App Service)

Azure Functions is a serverless solution that allows you to write less code, maintain less infrastructure, and save on costs. Instead of worrying about deploying and maintaining servers, the cloud infrastructure provides all the up-to-date resources needed to keep your applications running. According to <https://learn.microsoft.com/en-us/azure/azure-functions/functions-overview?pivots=programming-language-csharp>

upvoted 1 times

🗨️ 👤 **sultanmr123** 12 months ago

B. WebJobs

WebJobs is the most appropriate service for this scenario because:

The question specifically mentions using the WebJobs SDK to design a triggered App Service background task.

WebJobs is a feature of Azure App Service that enables you to run background tasks in the same instance as a web app.

The WebJobs SDK is designed specifically for running background tasks in Azure App Service, and it provides support for triggered jobs that can be automatically invoked when new data is received in a queue.

WebJobs can be configured to run continuously and respond to events, such as messages being added to a queue.

The scenario describes a task that automatically invokes a function when new data is received in a queue, which aligns perfectly with the capabilities of WebJobs and the WebJobs SDK.

While Azure Functions (option D) could also potentially be used for this scenario, the question explicitly mentions using the WebJobs SDK, making WebJobs the most appropriate choice for this specific use case.

upvoted 3 times

🗨️ 👤 **L\_\_\_L** 1 year, 2 months ago

**Selected Answer: B**

Both WebJob and Function are acceptable, but the context mentions "use the WebJobs SDK to design". WebJobs SDK is a framework designed for WebJobs that simplifies writing code that responds to events in Azure services. I prefer B.WebJob

upvoted 1 times

🗨️ 👤 **man5484** 1 year, 2 months ago

**Selected Answer: B**

The scenario described suggests that you're using the WebJobs SDK to create a triggered background task that processes data from a queue. Therefore, the correct service to use for this purpose is:

B. WebJobs

WebJobs allow you to run scripts or programs as background processes alongside your web app in Azure App Service. They can be triggered by various sources, including queues, blobs, and timers. In this case, since you're triggering a function every time new data is received in a queue, WebJobs would be the appropriate service to configure and manage this process.

upvoted 1 times

🗨️ 👤 **Mohanraj\_SV5** 1 year, 6 months ago

**Selected Answer: B**

B - WebJobs is correct

upvoted 1 times

🗨️ 👤 **HBSHAJER** 1 year, 8 months ago

HE ASK ABOUT A SERVICE NOT A FEATURE SO IT IS D

upvoted 1 times

🗨️ 👤 **person12345** 1 year, 8 months ago

There is an issue with the question. There is no mention of the OS where the app of the question is running. In the documentation Microsoft specifies that: "WebJobs aren't supported for App Service on Linux yet"

<https://learn.microsoft.com/en-us/azure/app-service/webjobs-create>

For that reason I think that the answer D) Azure Functions should be a right answer too.

upvoted 2 times

Your company has an Azure subscription.

You need to deploy a number of Azure virtual machines to the subscription by using Azure Resource Manager (ARM) templates. The virtual machines will be included in a single availability set.

You need to ensure that the ARM template allows for as many virtual machines as possible to remain accessible in the event of fabric failure or maintenance.

Which of the following is the value that you should configure for the platformFaultDomainCount property?

- A. 10
- B. 30
- C. Min Value
- D. Max Value

**Suggested Answer: D**

The number of fault domains for managed availability sets varies by region - either two or three per region.

Reference:

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

Community vote distribution

D (100%)

  **jay158** Highly Voted 9 months, 1 week ago

Answer D Max Value. [ values could be 1,2,3 ]

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-manage-fault-domains>  
upvoted 24 times

  **argoth** 3 years, 11 months ago

Correct. Values can be 1,2 or 3

upvoted 6 times

  **Bere** Highly Voted 9 months, 1 week ago

As described here:

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-manage-fault-domains>

You can set the parameter --platform-fault-domain-count to 1, 2, or 3 (default of 3 if not specified).

And as described here:

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

Each virtual machine in your availability set is assigned an update domain and a fault domain by the underlying Azure platform. Each availability set can be configured with up to three fault domains and twenty update domains.

So answer is D Max Value

upvoted 7 times

  **Ajju5791** Most Recent 6 months ago

Selected Answer: D

az vmss create \

--resource-group myResourceGroup \

--name myScaleSet \

--orchestration-mode Flexible \

--image Ubuntu2204 \

--admin-username azureuser \

--platform-fault-domain-count 3\

--generate-ssh-keys

upvoted 2 times

  **Zezere** 7 months, 1 week ago

Why is that relevant for the az-204? Is it not only an infrastructure question?

upvoted 2 times

🗳️ 👤 **Nhantran20000502** 8 months, 1 week ago

Selected Answer: D

Setting the platformFaultDomainCount property to the maximum value ensures that the virtual machines are distributed across the maximum number of fault domains. This maximizes the availability of the VMs in the event of a fabric failure or maintenance

upvoted 1 times

🗳️ 👤 **forgetfulalligator** 9 months, 1 week ago

Selected Answer: D

you should configure the platformFaultDomainCount property to its maximum value, which is 3 for most regions in Azure. The platformFaultDomainCount property specifies the number of fault domains to use for the availability set. By configuring it to its maximum value, the virtual machines will be distributed across multiple fault domains, which will help ensure that as many virtual machines as possible remain accessible in the event of a failure or maintenance event.

upvoted 3 times

🗳️ 👤 **man5484** 9 months, 1 week ago

Selected Answer: D

The platformFaultDomainCount property in an Azure Resource Manager (ARM) template specifies the number of fault domains for the Azure virtual machines within an availability set. Fault domains are distinct groups of hardware within a datacenter. When you deploy virtual machines within an availability set, Azure ensures that each VM is placed into a different fault domain to minimize the impact of hardware failures or maintenance events.

The maximum value allowed for platformFaultDomainCount depends on the region where you're deploying your resources. It varies across Azure regions. Typically, the maximum value is either 2 or 3. Therefore, to ensure that the ARM template allows for as many virtual machines as possible to remain accessible in the event of fabric failure or maintenance, you should configure

upvoted 1 times

🗳️ 👤 **Mohanraj\_SV5** 1 year, 6 months ago

Selected Answer: D

D. Max Value (As per the requirement)

upvoted 1 times

🗳️ 👤 **TheMCT** 1 year, 9 months ago

Selected Answer: D

Azure allows you to specify a platform fault domain count of up to 3, which means that there can be up to three fault domains in an availability set. This setting helps distribute the virtual machines across multiple fault domains, reducing the risk of simultaneous failure due to hardware issues. Therefore, you should set it to the maximum value of 3 for optimal fault tolerance.

upvoted 1 times

🗳️ 👤 **SmartTechieGrv** 2 years, 2 months ago

Answer is D

upvoted 1 times

🗳️ 👤 **p2006** 2 years, 4 months ago

<https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-manage-fault-domains#azure-cli>

upvoted 1 times

🗳️ 👤 **uffuchsi** 2 years, 4 months ago

For domain count, use D. Max value (Available options are 1, 2, 3)

upvoted 1 times

🗳️ 👤 **Santmax** 2 years, 5 months ago

D Max value

upvoted 1 times

🗳️ 👤 **kkstays** 2 years, 6 months ago

Selected Answer: D

Max value is Correct



upvoted 1 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

Selected Answer: D

Answer D Max Value

upvoted 1 times

  **braj1234** 2 years, 10 months ago

Answer is D

upvoted 1 times

  **ranjitklive** 3 years, 11 months ago

Each virtual machine in your availability set is assigned an update domain and a fault domain by the underlying Azure platform. Each availability set can be configured with up to three fault domains and twenty update domains. Update domains indicate groups of virtual machines and underlying physical hardware that can be rebooted at the same time.

upvoted 5 times

Your company has an Azure subscription.

You need to deploy a number of Azure virtual machines to the subscription by using Azure Resource Manager (ARM) templates. The virtual machines will be included in a single availability set.

You need to ensure that the ARM template allows for as many virtual machines as possible to remain accessible in the event of fabric failure or maintenance.

Which of the following is the value that you should configure for the platformUpdateDomainCount property?

- A. 10
- B. 20
- C. 30
- D. 40

**Suggested Answer: D**

Each virtual machine in your availability set is assigned an update domain and a fault domain by the underlying Azure platform. For a given availability set, five non-user-configurable update domains are assigned by default (Resource Manager deployments can then be increased to provide up to 20 update domains) to indicate groups of virtual machines and underlying physical hardware that can be rebooted at the same time.

Reference:

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

Community vote distribution

B (100%)

 **jay158** Highly Voted 9 months, 1 week ago

Answer is B 20

Each availability set can be configured with up to three fault domains and twenty update domains.

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

upvoted 58 times

 **ndh103** Highly Voted 3 years, 12 months ago

Should be 20

upvoted 11 times

 **Thulas** Most Recent 3 weeks, 6 days ago

**Selected Answer: A**

The correct answer is: A. 10

Explanation:

In an Azure availability set, the platformUpdateDomainCount property specifies how many update domains Azure can use for planned maintenance (e.g., OS or host updates).

Update domains (UDs) ensure that not all virtual machines (VMs) in the availability set are updated or rebooted at the same time during planned maintenance.

Azure allocates VMs across update domains automatically, up to a maximum of 20, but you can only configure up to 10 using ARM templates.

The default and maximum configurable value for platformUpdateDomainCount in an ARM template is 10.

Setting this value to 10 provides maximum distribution across update domains, improving the resiliency of your VMs during planned maintenance.


upvoted 1 times

 **Nhantran20000502** 8 months, 1 week ago

**Selected Answer: B**

In Azure, an availability set can be configured with up to 20 update domains. Setting the platformUpdateDomainCount property to 20 ensures that your virtual machines are distributed across the maximum number of update domains, which helps maintain accessibility during fabric failures or maintenance.

upvoted 2 times

 **Bogdan75** 9 months, 1 week ago

**Selected Answer: B**

"Update domains indicate groups of virtual machines and underlying physical hardware that can be rebooted at the same time. When more than five virtual machines are configured within a single availability set with five update domains, the sixth virtual machine is placed into the same update domain as the first virtual machine, the seventh in the same update domain as the second virtual machine, and so on."

Given that we want to maximize the number of VMs available at all times, we want to maximize the number of update domains – e.g. a single update domain would basically mean no redundancy at all. Having said that, there is a constraint on the maximum number of update domains:

"Each availability set can be configured with up to three fault domains and twenty update domains."

Source: <https://docs.microsoft.com/en-us/azure/virtual-machines/availability-set-overview#how-do-availability-sets-work>

upvoted 2 times


 **forgetfulalligator** 9 months, 1 week ago

**Selected Answer: B**

So I know the answer is B for sure, but still trying to wrap my head around the relationship between fault domain, update domain, and availability sets. Someone, please correct me, but from my understanding. Let's say I add 60 virtual machines to an availability set. Azure would distribute the VM across the 20 domain update domains, with each update domain containing three virtual machines. These Fault Domains can hold multiple Update Domain Groups. For Example, Fault Domain 1 could hold Update Domain Group 1 - 6. So if I were to do maintenance to Update domain 1, my clients would still have available for update domains 2 - 6. However, if Fault Domain 1 Goes down, then all Update Domains 1 -6 will also go down.


<https://learn.microsoft.com/bs-latn-ba/azure/virtual-machines/availability-set-overview#how-do-availability-sets-work>

upvoted 1 times

 **NK203** 2 years, 4 months ago

The Azure platform automatically distributes your VMs across these Fault domain.<https://learn.microsoft.com/en-us/training/modules/provision-virtual-machines-azure/3-azure-virtual-machine-availability-options?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.az-204-implement-iaas-solutions>

upvoted 1 times

 **Anjan1980** 9 months, 1 week ago

**Selected Answer: B**

The correct answer is B. 20.

The platformUpdateDomainCount property in Azure Resource Manager (ARM) templates specifies the number of update domains that should be created for the availability set. The maximum value for this property is 20.

By configuring the platformUpdateDomainCount property to the maximum value, you can help to ensure that your virtual machines are distributed across as many update domains as possible, which will help to protect them from maintenance operations that require updating virtual machines in a single update domain.

upvoted 1 times

 **man5484** 9 months, 1 week ago


**Selected Answer: B**

The `platformUpdateDomainCount` property in an Azure Resource Manager (ARM) template specifies the number of update domains for the Azure virtual machines within an availability set. Update domains are groups of VMs that are updated and rebooted together during planned maintenance events to minimize downtime for applications.

The maximum value allowed for `platformUpdateDomainCount` is 20. Therefore, to ensure that the ARM template allows for as many virtual machines as possible to remain accessible in the event of fabric failure or maintenance, you should configure:

B. 20

upvoted 2 times

 **Moe\_121** 9 months, 1 week ago

**Selected Answer: B**

Each virtual machine in your availability set is assigned an update domain and a fault domain by the underlying Azure platform. Each availability set can be configured with up to 3 fault domains and 20 update domains.

upvoted 1 times



🗨️ 👤 **Amir1909** 1 year, 4 months ago

B is correct

upvoted 1 times

🗨️ 👤 **Xoxoo** 1 year, 4 months ago

**Selected Answer: B**

The maximum value you can configure for platformUpdateDomainCount is 20. This means that virtual machines within the availability set will be spread across 20 update domains. By setting this value to the maximum, you ensure that the virtual machines are distributed optimally for availability during maintenance or update operations.

upvoted 3 times

🗨️ 👤 **Rouh** 1 year, 4 months ago

**Selected Answer: B**

Answer is B 20

upvoted 1 times

🗨️ 👤 **Mohanraj\_SV5** 1 year, 6 months ago

**Selected Answer: B**

The upper limit of platformUpdateDomainCount is 20

upvoted 2 times

🗨️ 👤 **Tarun\_A** 1 year, 7 months ago

**Selected Answer: B**

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

upvoted 2 times

🗨️ 👤 **regux** 1 year, 8 months ago

**Selected Answer: B**

B 20 is the correct one.

upvoted 1 times

🗨️ 👤 **Kopparthi** 1 year, 10 months ago

A higher value would generally be preferable to improve availability. Among the options, D. 40 would provide the highest degree of isolation and availability during maintenance or updates. Therefore, you should configure the platformUpdateDomainCount property with a value of 40

upvoted 1 times

🗨️ 👤 **Kopparthi** 1 year, 10 months ago

A higher value would generally be preferable to improve availability. Among the options, D. 40 would provide the highest degree of isolation and availability during maintenance or updates. Therefore, you should configure the platformUpdateDomainCount property with a value of 40

upvoted 1 times

DRAG DROP -

You are creating an Azure Cosmos DB account that makes use of the SQL API. Data will be added to the account every day by a web application.

You need to ensure that an email notification is sent when information is received from IoT devices, and that compute cost is reduced.

You decide to deploy a function app.

Which of the following should you configure the function app to use? Answer by dragging the correct options from the list to the answer area.

Select and Place:

## Options

Azure Cosmos DB  
connector

SendGrid action

Consumption plan

Azure Event Hubs  
binding

SendGrid binding

## Answer

# Options

# Answer

Azure Cosmos DB  
connector

SendGrid action

Azure Event Hubs  
binding

Consumption plan

SendGrid binding

Suggested Answer:

 **[Removed]** Highly Voted 3 years, 2 months ago

Answer is correct: Consumption plan will reduce the cost and SendGrid is used to send emails from azure functions.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-sendgrid?tabs=in-process%2Cfunctionsv2&pivots=programming-language-csharp>

upvoted 22 times

 **man5484** Highly Voted 1 year, 2 months ago

To meet the requirements of sending an email notification when information is received from IoT devices and reducing compute costs, you should configure the function app with the following options:

Azure Event Hubs binding: This will allow the function app to listen to events from IoT devices.

SendGrid action: This will enable the function app to send email notifications.

Consumption plan: This plan offers automatic scaling and only charges for the resources consumed by the function app, which helps in reducing compute costs.

upvoted 6 times

 **shane0909** Most Recent 6 days, 22 hours ago

Consumption plan: This helps in reducing compute costs as the consumption plan is a pay-per-use model, meaning you only pay for the time the function runs.

SendGrid action: This should be used to send email notifications in the function app when an event happens, such as when IoT data is received.

Azure Cosmos DB connector: Not focused on sending email notifications or reducing compute cost.

SendGrid binding: Although it could send emails, SendGrid action is more specific to triggering an email as part of an action in response to an event.

Azure Event Hubs binding: Useful for IoT data streaming but not directly related to the specific task of sending email notifications and reducing compute costs.

upvoted 1 times

🗨️ 👤 **MortisG** 3 weeks, 4 days ago

Answer is correct. Consumption Plan and SendGrid Binding meet the requirements and both apply to Azure Functions (all others don't).

Cosmos DB Connector is a Logic Apps term (not functions) so it doesn't apply to this scenario.

upvoted 2 times

🗨️ 👤 **ankita0910** 4 months, 1 week ago

To meet your requirements, you should select the following two options:

Consumption Plan: This will help reduce compute costs by only charging for the execution time of your functions.

SendGrid Binding: Use this binding to send email notifications when new data is received from IoT devices.

upvoted 1 times

🗨️ 👤 **fiercefreesom** 9 months ago

In a real application of this situation, you would need some type of trigger coming from CosmosDB in order to notify the app to send the email.

CosmosDB Connectors are only for Logic Apps, and other CosmosDB triggers aren't part of the options. That's what makes this problem tricky and unrealistic.

upvoted 3 times

🗨️ 👤 **forgetfulalligator** 9 months, 1 week ago

Hmm super tough problem , I can completely understand why people would choose Azure Event Hubs cause you could receive and process messages from IoT devices, and can trigger the function app to perform certain actions when new data is received. For example, you could use an Event Hub binding to send an email notification when new data is received from an IoT device. But it's most likely SendGrid binding since it could technically do it (I guess) but also it would keep computing costs low since we are processing a low-volume and low-latency task and it is specific to sending emails. Azure Event Hubs binding is usually for processing loads of events so not the best use case here.

upvoted 4 times

🗨️ 👤 **AryaNiaz** 10 months, 2 weeks ago

Both options are indeed correct; however, regarding the function trigger, shouldn't it be the Azure Cosmos DB connector serving as the trigger?

upvoted 1 times

🗨️ 👤 **Christian\_garcia\_martin** 10 months, 2 weeks ago

The answer is correct ,i would add a cosmos connector to interact with cosmos DB .

upvoted 1 times

🗨️ 👤 **Christian\_garcia\_martin** 10 months, 2 weeks ago

nevermind ,DB connector should be for logic app no for function so answer is correct as given

upvoted 1 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

Seems to be correct answer after my investigation.

upvoted 2 times

🗨️ 👤 **whitefancy** 1 year, 6 months ago

what is different between sendgrid binding and sendgrid action?

upvoted 2 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

If your primary goal is to send emails as part of a broader workflow, you might find SendGrid binding more convenient. If you need more control or want to customize the email-sending logic, using SendGrid as a direct action in your function code may be preferable.

upvoted 3 times

🗨️ 👤 **Aneeqa** 1 year, 9 months ago

SendGrid action: This option is used to send email messages using SendGrid. However, it requires you to write code to send the email messages. By using the SendGrid binding, you can easily send email notifications without writing any code.

Azure Event Hubs binding: This option is used to receive and process events from Azure Event Hubs. It is not relevant for sending email notifications.

upvoted 2 times

🗨️ 👤 **zreaf29** 1 year, 11 months ago

Anyone can explain me what is different between SendGrid action and SendGrid binding??

upvoted 1 times

🗨️ 👤 **Molnmannen1337** 1 year, 11 months ago

Actions are used in logic apps, bindings in functions. That is my guess.

upvoted 4 times

🗨️ 👤 **zreaf29** 9 months, 1 week ago

I searched this question in Google Bard. And I got the answer.

Both two are functions. But the difference is SendGrid Action is used to send an email synchronously, while SendGrid binding is used to send an email asynchronously.

It means that SendGrid Action will block the execution of the function app until the email has been sent. This can lead to increased compute costs, especially if you are sending a large number of emails. On the other way, SendGrid binding will not block while the email is being sent. This can help to reduce compute costs, as the function app can continue to process other requests while the email is being sent.

Question is "You need to ensure that compute cost is reduced". So I think, Answer is SendGrid binding which is good for reducing costs.

upvoted 5 times

🗨️ 👤 **HQ1227** 1 year, 12 months ago

what is different between sendgrid binding and sendgrid action?

upvoted 3 times

🗨️ 👤 **LuchoGonzalez** 2 years ago

Based on the options provided, the recommended configuration for the function app would be to use the "SendGrid action" for sending email notifications and the "consumption plan" to reduce compute costs. Check it out this pls

upvoted 2 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

Consumption Plan (Keep costs low) & Send Grid Bunding (Send Email). I think it could also be a Function App and Event Hub, but initial answer is more accurate.

upvoted 2 times

🗨️ 👤 **Alluru** 2 years, 8 months ago

Configure the function app to use consumption plan and use sendgrid bindings to send emails.

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-sendgrid?tabs=in-process%2Cfunctionsv2&pivots=programming-language-csharp#example>

upvoted 1 times

This question requires that you evaluate the underlined text to determine if it is correct.

You company has an on-premises deployment of MongoDB, and an Azure Cosmos DB account that makes use of the MongoDB API.

You need to devise a strategy to migrate MongoDB to the Azure Cosmos DB account.

You include the Data Management Gateway tool in your migration strategy.

Instructions: Review the underlined text. If it makes the statement correct, select 'No change required.' If the statement is incorrect, select the answer choice that makes the statement correct.

- A. No change required
- B. mongorestore
- C. Azure Storage Explorer
- D. AzCopy

**Suggested Answer: B**

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/mongodb-migrate> <https://docs.mongodb.com/manual/reference/program/mongorestore/>

Community vote distribution

B (84%)

A (16%)

 **jay158** Highly Voted 9 months, 1 week ago

Answer: B --mongorestore

<https://docs.microsoft.com/en-us/azure/cosmos-db/mongodb-pre-migration>

upvoted 35 times

 **jay158** 3 years, 11 months ago

Azure Database Migration Service [ is for online Migration] but it is not a choice

<https://docs.microsoft.com/en-us/azure/dms/tutorial-mongodb-cosmos-db-online>

upvoted 4 times

 **AzureLearning** 3 years, 5 months ago

The Data management gateway is a client agent that you must install in your on-premises environment to copy data between cloud and on-premises data stores.

<https://docs.microsoft.com/en-us/azure/data-factory/v1/data-factory-data-management-gateway>

upvoted 2 times

 **forgetfulalligator** Highly Voted 2 years, 4 months ago

**Selected Answer: B**

mongorestore can be used to restore a MongoDB backup to the Azure Cosmos DB account using the MongoDB API.

upvoted 6 times

 **Nhantran20000502** Most Recent 8 months, 1 week ago

**Selected Answer: B**

The underlined text is incorrect. The Data Management Gateway tool is not typically used for migrating MongoDB to Azure Cosmos DB.

The correct answer is B. mongorestore.

To migrate MongoDB to Azure Cosmos DB, you can use MongoDB native tools like mongodump and mongorestore. These tools are designed to handle the export and import of MongoDB data efficiently.

upvoted 1 times

 **Mol\_** 9 months, 1 week ago

**Selected Answer: B**

Got this on the Microsoft documentation: "The MongoDB native tools are a set of binaries that facilitate data manipulation on an existing MongoDB instance. Since Azure Cosmos DB exposes an API for MongoDB, the MongoDB native tools are able to insert data into Azure Cosmos DB. The focus of this doc is on migrating data out of a MongoDB instance using mongoexport/mongoimport or mongodump/mongorestore. Since the native tools

connect to MongoDB using connection strings, you can run the tools anywhere, however we recommend running these tools within the same network as the MongoDB instance to avoid firewall issues."

upvoted 1 times

🗨️ 👤 **man5484** 9 months, 1 week ago

**Selected Answer: B**

Data Management Gateway: This tool is used for connecting on-premises data sources (like SQL Server) to Azure Analysis Services or Power BI for data visualization and reporting. It's not designed for direct database migrations.

Correct Answer:

B. mongorestore

The mongorestore command-line tool is specifically designed for restoring data backups created with mongodump into a MongoDB instance. This makes it a suitable option for migrating data from your on-premises MongoDB deployment to the Azure Cosmos DB account using the MongoDB API.

upvoted 3 times

🗨️ 👤 **harbox** 9 months, 1 week ago

**Selected Answer: B**

<https://docs.microsoft.com/en-us/azure/cosmos-db/mongodb-migrate> <https://docs.mongodb.com/manual/reference/program/mongorestore/>

The underlined text states that the Data Management Gateway tool is included in the migration strategy for migrating MongoDB to the Azure Cosmos DB account that uses the MongoDB API. This statement is incorrect.

The Data Management Gateway tool is used to connect on-premises data sources to Azure services, including Power BI, PowerApps, and Microsoft Flow. It provides a secure and reliable connection between on-premises data sources and cloud-based services.

However, in the context of migrating MongoDB to Azure Cosmos DB, the Data Management Gateway tool is not relevant. Instead, you should use one of the available migration tools such as the Azure Database Migration Service or the mongoimport/mongodump utilities to migrate your data.

Therefore, the correct answer is B. mongorestore.

upvoted 2 times

🗨️ 👤 **Dianahu** 1 year, 11 months ago

No change required => <https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-on-premises-mongodb-connector> and <https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-move-data-between-onprem-and-cloud> all the links supporting option with mongorestore do not mention data management gateway. For data management gateway you only need to specify schema and the rest is done by DMG and its pipeline with a copy activity.

upvoted 1 times

🗨️ 👤 **Alluru** 2 years, 3 months ago

Data Management Gateway is not supported for Azure Cosmos DB - Mongo API

<https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-data-movement-activities#supported-data-stores-and-formats>

So you can use native tools to perform an offline (one-time) migration of databases from an on-premises or cloud instance of MongoDB to Azure Cosmos DB's API for MongoDB.

<https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/tutorial-mongotools-cosmos-db#overview-of-data-migration-using-mongodb-native-tools>

upvoted 1 times

🗨️ 👤 **JhonR** 2 years, 2 months ago

The question doesn't mention DataFactory, which is where your first link points.

Here we can see DMG is listed among the migration tools: <https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/pre-migration-steps>.

So the answer is A.

upvoted 1 times

🗨️ 👤 **wrak** 2 years, 4 months ago

When moving the whole DB, we are using mongorestore: <https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/tutorial-mongotools-cosmos-db>

upvoted 1 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

A - No change required

upvoted 1 times

🗨️ 👤 **fkaran** 2 years, 4 months ago

To migrate MongoDB to an Azure Cosmos DB account that uses the MongoDB API, you can use mongodump and mongorestore tools that are included with MongoDB.

upvoted 3 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

I think "MongoDB API" is a key

upvoted 1 times

🗨️ 👤 **guschess** 2 years, 5 months ago

**Selected Answer: A**

"include data management gateway"

<https://azure.microsoft.com/en-us/blog/data-management-gateway-high-availability-and-scalability-preview/>

upvoted 1 times

🗨️ 👤 **armpro** 2 years, 5 months ago

**Selected Answer: A**

Answer: A

In Data Management Gateway i.e., Data factory once source and sink are setup, there is no need for mongorestore to run, running the pipeline will transfer data.

upvoted 2 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

Are you calling setting up a gateway and using a client agent a "tool"?

upvoted 1 times

🗨️ 👤 **lormar72** 2 years, 6 months ago

i think is A the DMG do the job

upvoted 1 times

🗨️ 👤 **kkop** 2 years, 7 months ago

**Selected Answer: B**

mongorestore imports your BSON file dump into Azure Cosmos DB.

Here, you can enhance the level of your knowledge about MongoDB migrations:

<https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/tutorial-mongotools-cosmos-db>

upvoted 4 times

🗨️ 👤 **Alluru** 2 years, 8 months ago

For datastores with <10 GB monogoDB native tool is used for migration from on-premises data store to cloud data store.

<https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/pre-migration-steps#pre-migration-logistics-planning>

upvoted 1 times

🗨️ 👤 **Alluru** 2 years, 8 months ago

I agree with Kozchris, the answer is B - mongorestore. DMG can not have Cosmos DB - Mongo as a supported sink (target).

<https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-data-movement-activities#supported-data-stores-and-formats>

upvoted 1 times



You are developing an e-Commerce Web App.

You want to use Azure Key Vault to ensure that sign-ins to the e-Commerce Web App are secured by using Azure App Service authentication and Azure Active Directory (AAD).

What should you do on the e-Commerce Web App?


- A. Run the `az keyvault secret` command.
- B. Enable Azure AD Connect.
- C. Enable Managed Service Identity (MSI).
- D. Create an Azure AD service principal.

**Suggested Answer: C**

A managed identity from Azure Active Directory allows your app to easily access other AAD-protected resources such as Azure Key Vault. Reference:

<https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity> <https://docs.microsoft.com/en-us/samples/azure-samples/app-service-msi-keyvault-dotnet/keyvault-msi-appservice-sample/>

Community vote distribution



**forgetfulalligator** Highly Voted 9 months, 1 week ago

**Selected Answer: C**

Managed Service Identity (MSI) is a feature in Azure that allows you to securely authenticate an Azure service to other Azure services without having to manage credentials. By enabling MSI on the Azure App Service hosting the e-Commerce Web App, you can create a trust relationship between the App Service and Azure Key Vault. This allows the e-Commerce Web App to authenticate with Azure Active Directory (AAD) and securely retrieve secrets from the Key Vault.

upvoted 9 times

**Nhantran20000502** Most Recent 8 months, 1 week ago

**Selected Answer: C**

Enabling Managed Service Identity (MSI) allows your e-Commerce Web App to securely access Azure Key Vault without needing to manage credentials. MSI provides an automatically managed identity in Azure AD, which can be used to authenticate to any service that supports Azure AD authentication, including Azure Key Vault.

upvoted 1 times

**junkz** 9 months, 1 week ago

since the statement says "secured by using Azure App Service authentication and Azure Active Directory (AAD)."

it stand to reason the nswer would be D. when configuring app service authentication with Microsoft identity provider, an app registration is required. which, under the hood, creates an app service principal (the ObjectID from the app registration, not to be confused with the applicationID of the app registration). so although the wording of the D option is missfortunate, it's still true.

upvoted 2 times

**SirMarv** 1 year, 12 months ago

**Selected Answer: C**

C is correct

upvoted 1 times

**uffuchsi** 2 years, 4 months ago

C - Enable Managed Service Identity

upvoted 1 times

**alexein74** 2 years, 5 months ago

D. Create an Azure AD service principal.

To secure sign-ins to the e-Commerce Web App by using Azure App Service authentication and Azure Active Directory (AAD), you should create an Azure AD service principal. A service principal is a security identity that you can use to authenticate and authorize your app to access Azure resources. Once you have created the service principal, you can use it to authenticate to Azure Key Vault and access the secrets that you store there.

B. Enable Azure AD Connect is a method to Synchronize on-premises directories and enable single sign-on and it's not related to this question.  
A. Run the az keyvault secret command is a command line to manage secrets in keyvault, it's not related to the question.  
C. Enable Managed Service Identity (MSI) is a feature that enables an app to authenticate to Azure services using its managed identity, which is automatically managed by Azure. This feature can be used to authenticate to Azure Key Vault and access the secrets that you store there, but it's not necessary to use it in this scenario.

upvoted 3 times

🗳️ 👤 **Mattt** 2 years, 3 months ago

Option D is incorrect because creating an Azure AD service principal is used to allow applications to authenticate and access resources in Azure, but it is not required to secure sign-ins to an e-Commerce Web App with Azure App Service authentication and Azure AD using Azure Key Vault.

upvoted 1 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: C**

C. Enable Managed Service Identity (MSI).

upvoted 2 times

🗳️ 👤 **KingChuang** 2 years, 6 months ago

While developers can securely store the secrets in Azure Key Vault, services need a way to access Azure Key Vault. Managed identities provide an automatically managed identity in Azure Active Directory for applications to use when connecting to resources that support Azure Active Directory (Azure AD) authentication. Applications can use managed identities to obtain Azure AD tokens without having to manage any credentials.

Ref: <https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>

upvoted 1 times

🗳️ 👤 **ericci** 2 years, 7 months ago

**Selected Answer: A**

I think the right answer is A: <https://learn.microsoft.com/en-us/azure/data-factory/v1/data-factory-on-premises-mongodb-connector>

upvoted 1 times

🗳️ 👤 **sca88** 2 years, 9 months ago

C is better than D, because thanks to ManagedIdentity, your code can forget to store keys, so is better solution than Service Principal

upvoted 2 times

🗳️ 👤 **PhilLI** 3 years, 5 months ago

**Selected Answer: C**

"Azure AD service principals" are created for Apps registered in Azure AD.

Whereas WebApps, VMs, hubs etc all can get a Managed Identity.

upvoted 4 times

🗳️ 👤 **xahah22222** 3 years, 8 months ago

I dont know. C&D seem to be the same thing. I would take the create principle option as I am not even sure if you can enable or disable MSI on a top level.

upvoted 1 times

🗳️ 👤 **Sukon\_Desknot** 3 years, 10 months ago

This is a bit tricky checkout

<https://docs.microsoft.com/en-us/azure/active-directory/develop/app-objects-and-service-principals#service-principal-object>

upvoted 4 times

This question requires that you evaluate the underlined text to determine if it is correct.

Your Azure Active Directory Azure (Azure AD) tenant has an Azure subscription linked to it.

Your developer has created a mobile application that obtains Azure AD access tokens using the OAuth 2 implicit grant type.

The mobile application must be registered in Azure AD.

You require a redirect URI from the developer for registration purposes.

Instructions: Review the underlined text. If it makes the statement correct, select 'No change is needed.' If the statement is incorrect, select the answer choice that makes the statement correct.

A. No change required.

B. a secret

C. a login hint

D. a client ID

#### Suggested Answer: A

For Native Applications you need to provide a Redirect URI, which Azure AD will use to return token responses.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/v1-protocols-oauth-code>

Community vote distribution



A (100%)

 **jvyas**  9 months, 1 week ago

You don't need client id to register an application in Azure AD. You just need redirect URI. Upon registration, the portal will give client id and tenant id. Both of them must be included in the app configuration json file, so given answer is correct as far as app registration is concerned.

<https://docs.microsoft.com/en-us/learn/modules/secure-app-with-oidc-and-azure-ad/4-exercise-create-aad-register-app>

upvoted 35 times

 **qwerty112233**  9 months, 1 week ago

"This question requires that you evaluate the underlined text to determine if it is correct." - where is the underlined text to determine? i did not get it  
upvoted 13 times

 **JustHereToLearn** 2 years, 5 months ago

I believe this is supposed to be the underlined text on which the answer seems correct "You require a redirect URI from the developer for registration purposes."

upvoted 7 times

 **Mirdwal**  1 month, 3 weeks ago

**Selected Answer: A**

When a mobile app uses OAuth 2.0 implicit grant flow (or more commonly today, the authorization code flow with PKCE), Azure AD needs a redirect URI to:

upvoted 1 times

 **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: A**

When a user tries to authenticate and authorize an application to access their Azure AD resources, Azure AD sends an authentication token to the redirect URI that was specified during application registration. The application then extracts the token from the redirect URL and uses it to access Azure AD resources on behalf of the user.

upvoted 5 times

 **wrak** 2 years, 4 months ago

As part of the app registration process two things are required: Client ID and Redirect URI. So the statement in the question is correct - You require a redirect URI from the developer for registration purposes. Hence we don't have to make any changes - answer A is correct

upvoted 1 times

 **uffuchsi** 2 years, 4 months ago

Since the client Id so not something the developer would necessarily know, you would only need them to provide the redirect URI, so A

upvoted 1 times

🗨️ 👤 **guschess** 2 years, 5 months ago

**Selected Answer: A**

You require a redirect (URI|client id|other options) from the developer for registration purposes.

upvoted 4 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: A**

I think the answer is correct. ClientID is generated and provided to the developer by the admin after registration that is used for subsequent use. i.e after registration the clientid is generate.

"No change is required"

upvoted 2 times

🗨️ 👤 **upadhyayavi** 3 years, 2 months ago

You get the client ID(or app ID) after the app is registered.

upvoted 2 times

🗨️ 👤 **Prashant\_Salokhe** 3 years, 3 months ago

This type of questions are misleading. No underline text.

upvoted 3 times

🗨️ 👤 **SivajiTheBoss** 3 years, 3 months ago

100% Answer is Correct:

Redirect URI from the developer for registration purposes will solve the purpose as the mobile application registered in Azure AD and obtains Azure AD access tokens using the OAuth 2 implicit grant type.

upvoted 1 times

🗨️ 👤 **Mev4953** 3 years, 5 months ago

Which text is underlined text?

upvoted 4 times

🗨️ 👤 **huislaw** 3 years, 5 months ago

I don't see it too

upvoted 2 times

🗨️ 👤 **Laueri** 3 years, 5 months ago

I believe in this case it's "redirect URI" that should be underlined.

upvoted 6 times

🗨️ 👤 **Patchfox** 3 years, 6 months ago

I walked through the registration process of the enterprise app but the redirect URL is optional to continue . I'm sure that you need it obviously later. But the trick is here, all other answers not needed too, so you have to go with one and the best way is "No change required"

upvoted 4 times

🗨️ 👤 **Patchfox** 3 years, 6 months ago

I want to add here the info from the official documentation, I read. The different authentication types are really confusing. So the documantion say, redirect\_uri is required for native and mobile applications. So I ahve to say, the answer here is correct. No change required.

upvoted 3 times

🗨️ 👤 **MiraA** 3 years, 8 months ago

Register an application with the Microsoft identity platform - add a redirect URI:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-register-app#add-a-redirect-uri>

upvoted 2 times

🗨️ 👤 **MiraA** 3 years, 8 months ago

"redirect\_uri" ... The redirect\_uri of your app, where authentication responses can be sent and received by your app. It must exactly match one of the redirect\_uris you registered in the portal, except it must be url encoded. For native & mobile apps, you should use one of the recommended values - <https://login.microsoftonline.com/common/oauth2/nativeclient> (for apps using embedded browsers) or <http://localhost> (for apps that use system browsers).



<https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-auth-code-flow>

upvoted 2 times

🗨️ 👤 **Omar102** 3 years, 9 months ago

Client ID is also a required.

upvoted 5 times

  **DV007** 3 years, 5 months ago

The question is about which data is required \*from the developer\*. The Client ID (and client secret) are provided TO the developer by you as admin. The thing the admin needs to know from the developer is the redirect URL, because that is what the developer knows for his application

upvoted 8 times

  **ucsdmiami2020** 3 years, 6 months ago

I agreed the answer is Client ID per the Microsoft docs reference URL <https://docs.microsoft.com/en-us/learn/modules/secure-app-with-oidc-and-azure-ad/4-exercise-create-aad-register-app>

"Redirect URI (optional) - We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios."

On the other hand,

"Hoover over the Application (client) ID value and copy to a secure location. Save the values for later. You'll need these to complete the configuration file"

upvoted 1 times

You are creating an Azure key vault using PowerShell. Objects deleted from the key vault must be kept for a set period of 90 days. Which two of the following parameters must be used in conjunction to meet the requirement? (Choose two.)

- A. EnabledForDeployment
- B. EnablePurgeProtection
- C. EnabledForTemplateDeployment
- D. EnableSoftDelete

**Suggested Answer: BD**

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.keyvault/new-azurermkeyvault> <https://docs.microsoft.com/en-us/azure/key-vault/key-vault-ovw-soft-delete>

Community vote distribution

BD (100%)


 **forgetfulalligator** Highly Voted 9 months, 1 week ago

**Selected Answer: BD**

The EnablePurgeProtection parameter prevents the deletion of key vault objects by enabling purge protection. This parameter ensures that objects deleted from the key vault cannot be permanently deleted until the purge protection period has expired.

The EnableSoftDelete parameter allows the key vault to retain deleted objects for a specified duration. This parameter enables soft delete, which allows the key vault to retain deleted objects for a specified period of time (90 days in this case) before they are permanently deleted.

upvoted 16 times

 **SagarUnhale** Most Recent 1 month, 2 weeks ago

**Selected Answer: BD**

Why not other options?

A. EnabledForDeployment

Used to allow virtual machines to retrieve secrets during deployment, not related to deletion or retention.

C. EnabledForTemplateDeployment

Enables access to secrets during Azure Resource Manager (ARM) template deployments, unrelated to retention policies.

Correct Answer

B. EnablePurgeProtection

D. EnableSoftDelete

upvoted 1 times

 **TheFivePips** 10 months, 2 weeks ago

**Selected Answer: BD**

Soft Delete is necessary to create a recoverable state for deleted objects. Without it, objects are immediately and permanently deleted, and Purge Protection has nothing to protect.

Purge Protection enhances the security provided by Soft Delete by ensuring that once an object is soft-deleted, it cannot be purged until after the retention period.

upvoted 1 times

 **vincent\_huang** 12 months ago

**Selected Answer: BD**

BD is correct.

upvoted 1 times

 **Archana\_G** 1 year, 1 month ago

**Selected Answer: BD**

BD is correct.

upvoted 1 times

🗄️ 👤 **viking2** 2 years ago

**Selected Answer: BD**

BD is correct. You'll need to enable soft delete, and then purge protection to make sure that soft-deleted objects are not purged early.

upvoted 4 times

🗄️ 👤 **UkadikeToch** 2 years ago

BD

EnabledForDeployment is a property of Azure Key Vault that allows you to control whether the Key Vault can be deployed as part of an Azure Resource Manager (ARM) template deployment.

When EnabledForDeployment is set to true, it means that the Key Vault can be included in ARM templates and deployed using Azure Resource Manager.

EnablePurgeProtection is a property of Azure Key Vault's soft delete feature.

When EnablePurgeProtection is set to true, it ensures that deleted Key Vault entities (such as secrets, keys, or certificates) cannot be permanently deleted immediately.

EnabledForTemplateDeployment is a property of Azure Key Vault that determines whether the Key Vault can be deployed as part of an Azure Resource Manager (ARM) template deployment.

EnableSoftDelete is a property of Azure Key Vault that enables the soft delete feature for the Key Vault.

Soft delete allows you to recover deleted entities (secrets, keys, or certificates) within a specified retention period, even after they have been deleted.

upvoted 3 times

🗄️ 👤 **AZAdam22** 2 years, 1 month ago

**Selected Answer: BD**

You need soft delete (D) to retain deleted objects for up to 90 days, and you need purge protection (B) to prevent deleted objects from being purged until 90 days have passed, as it says 90 "set" days.

upvoted 1 times

🗄️ 👤 **uffuchsi** 2 years, 4 months ago

B, and C. Purge defines how long a blob is 'alive' before it can be deleted. Soft delete allows you to recover mistakenly deleted blobs

upvoted 1 times

🗄️ 👤 **uffuchsi** 2 years, 4 months ago

Mis-type, meant to be B, and D

upvoted 1 times

🗄️ 👤 **alexein74** 2 years, 5 months ago

B. EnablePurgeProtection

D. EnableSoftDelete

To meet the requirement of keeping objects deleted from the key vault for 90 days, you need to use the EnablePurgeProtection and EnableSoftDelete parameters when creating the key vault using PowerShell.

EnablePurgeProtection prevents the key vault from being permanently deleted before the soft-delete retention period has elapsed.

EnableSoftDelete allows deleted vault and its contents to be retained and recoverable for the specified number of days.

upvoted 3 times

🗄️ 👤 **HafizSalmanMalik** 2 years, 6 months ago

Correct options are

B. EnablePurgeProtection

D. EnableSoftDelete

upvoted 3 times

🗄️ 👤 **Topic9** 2 years, 6 months ago

B & D.

upvoted 1 times

🗄️ 👤 **elequiel** 2 years, 7 months ago

**Selected Answer: BD**

Correct.

upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: BD**

correct - B and D

upvoted 1 times

🗨️ 👤 **Magis** 2 years, 7 months ago

**Selected Answer: BD**

Correct.

upvoted 2 times

🗨️ 👤 **helnik** 2 years, 8 months ago

correct b, d

upvoted 1 times



## HOTSPOT -

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

You want to implement multi-factor authentication by making use of a conditional access policy. The conditional access policy must be applied to all users when they access the Azure portal.

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

The screenshot shows a configuration window titled "Require MFA for Azure portal access". It has a dark header bar with a close button. Below the header, there are "Info" and "Delete" buttons. The main content area is divided into sections:

- Name:** A text box containing "MFA required for Azure portal access".
- Assignments:** A section with three expandable items:
  - Users and groups:** 0 users selected
  - Cloud apps:** 0 apps selected
  - Conditions:** 0 conditions selected
- Access controls:** A section with two expandable items:
  - Grant:** 0 controls selected
  - Session:** 0 controls selected

## Answer Area

Require MFA for Azure port... [icon] [X]

Info [icon] Delete [icon]

\* Name  
MFA required for Azure portal access

Assignments

- Users and groups [icon] 0 users selected >
- Cloud apps [icon] 0 apps selected >
- Conditions [icon] 0 conditions selected >

Access controls

- Grant [icon] 0 controls selected >
- Session [icon] 0 controls selected >

Suggested Answer:

Box 1:

The conditional access policy must be applied or assigned to Users and Groups.

Box 2:

The conditional access policy must be applied when users access the Azure portal, which is a cloud app. That is: Microsoft Azure Management

Box 3:

Access control must require multi-factor authentication when granting access.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/app-based-mfa>

**finnishr** Highly Voted 9 months, 1 week ago

The answer is correct! Please refer to:

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/howto-conditional-access-policy-azure-management>  
upvoted 15 times

**Ody** 6 months, 1 week ago

The answer is correct, but Cloud Apps has been changed to "Target Resources".  
upvoted 1 times

**Nestiei** Most Recent 9 months, 1 week ago

The answer is correct.

It's explicit here:

<https://learn.microsoft.com/en-us/azure/active-directory/conditional-access/howto-conditional-access-policy-all-users-mfa#create-a-conditional-access-policy>  
upvoted 4 times

**656823** 1 year, 10 months ago

Pardon but I must have forgotten, why not select Conditions?  
upvoted 2 times



**forgetfulalligator** 2 years, 4 months ago

Assignments: Include all users or a specific group that you want to apply the policy to.

Cloud apps or actions: Select "Microsoft Azure Management" or "Microsoft Azure portal".


Access controls: Configure the policy to require multi-factor authentication.

upvoted 2 times

  **uffuchsi** 2 years, 4 months ago

I originally thought Users and groups, Conditions, and Grant. As the removed suggests, the requirement (condition) for MFA not in the "conditions" tab, but in the "grants" thus the correct answer is not Conditions but instead Cloud apps

upvoted 1 times

  **OPT\_001122** 2 years, 7 months ago

The answer is correct

upvoted 1 times

You manage an Azure SQL database that allows for Azure AD authentication.

You need to make sure that database developers can connect to the SQL database via Microsoft SQL Server Management Studio (SSMS). You also need to make sure the developers use their on-premises Active Directory account for authentication. Your strategy should allow for authentication prompts to be kept to a minimum.

Which of the following should you implement?

- A. Azure AD token.
- B. Azure Multi-Factor authentication.
- C. Active Directory integrated authentication.
- D. OATH software tokens.

**Suggested Answer: C**

Azure AD can be the initial Azure AD managed domain. Azure AD can also be an on-premises Active Directory Domain Services that is federated with the Azure AD.

Using an Azure AD identity to connect using SSMS or SSDT

The following procedures show you how to connect to a SQL database with an Azure AD identity using SQL Server Management Studio or SQL Server Database

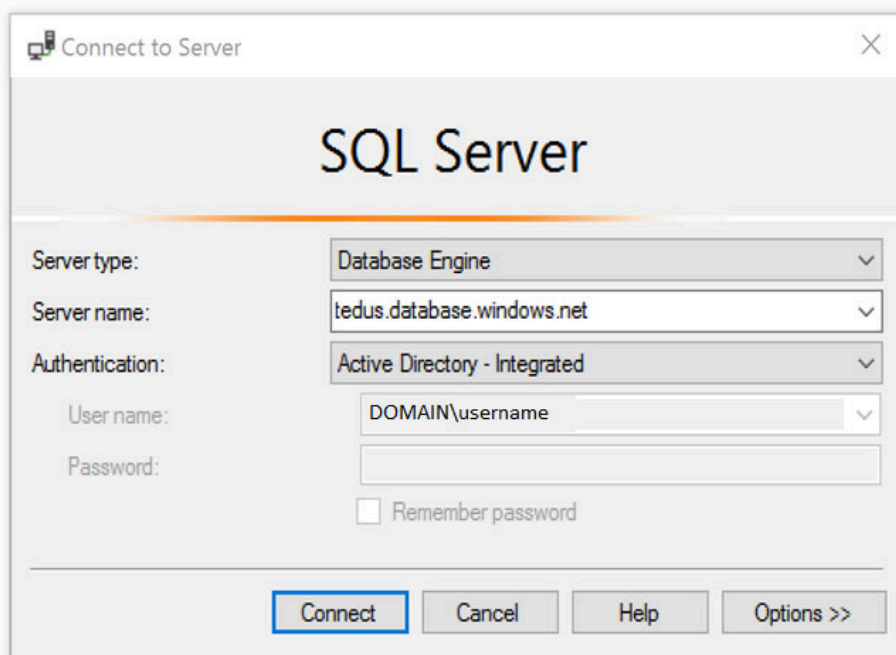
Tools.

Active Directory integrated authentication

Use this method if you are logged in to Windows using your Azure Active Directory credentials from a federated domain.

1. Start Management Studio or Data Tools and in the Connect to Server (or Connect to Database Engine) dialog box, in the Authentication box, select Active

Directory - Integrated. No password is needed or can be entered because your existing credentials will be presented for the connection.



2. Select the Options button, and on the Connection Properties page, in the Connect to database box, type the name of the user database you want to connect to.

(The AD domain name or tenant ID option is only supported for Universal with MFA connection options, otherwise it is greyed out.)

Community vote distribution

C (100%)

MrXBasit Highly Voted 3 years, 11 months ago

Correct

upvoted 26 times

forgetfulalligator Highly Voted 9 months, 1 week ago

**Selected Answer: C**

Active Directory integrated authentication is the correct option as it allows users to connect to the database using their Windows credentials, which are authenticated through their on-premises Active Directory. This option avoids the need for users to enter their credentials each time they connect to the database, reducing authentication prompts to a minimum.

upvoted 6 times

🗲️ 👤 **Gomza** **Most Recent** 1 month, 3 weeks ago

**Selected Answer: C**

Correct

upvoted 1 times

🗲️ 👤 **6d58417** 7 months ago

**Selected Answer: C**

<https://learn.microsoft.com/en-us/sql/connect/ado-net/sql/azure-active-directory-authentication?view=sql-server-ver16#using-integrated-authentication>

upvoted 1 times

🗲️ 👤 **harbox** 9 months, 1 week ago

**Selected Answer: C**

Azure AD can be the initial Azure AD managed domain.

Azure AD can also be an on-premises Active Directory Domain Services that is federated with the Azure AD.

Using an Azure AD identity to connect using SSMS or SSDT The following procedures show you how to connect to a SQL database with an Azure AD identity using SQL Server Management Studio or SQL Server Database Tools.

Active Directory integrated authentication Use this method if you are logged in to Windows using your Azure Active Directory credentials from a federated domain.

upvoted 1 times

🗲️ 👤 **UkadikeToch** 9 months, 1 week ago

Active Directory integrated authentication allows developers to use their on-premises AD account to authenticate and connect to the Azure SQL database.

With this method, SSMS uses the logged-in Windows credentials of the user to authenticate against the database.

By configuring SSO between on-premises AD and Azure AD, developers can achieve a seamless authentication experience without additional prompts.

upvoted 4 times

🗲️ 👤 **TusharChandak** 1 year, 11 months ago

**Selected Answer: C**

Correct

upvoted 1 times

🗲️ 👤 **uffuchsi** 2 years, 4 months ago

C - Active Directory integrated authentication

upvoted 1 times

🗲️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: C**

C. Active Directory integrated authentication.

upvoted 2 times

You are developing an application to transfer data between on-premises file servers and Azure Blob storage. The application stores keys, secrets, and certificates in Azure Key Vault and makes use of the Azure Key Vault APIs.

You want to configure the application to allow recovery of an accidental deletion of the key vault or key vault objects for 90 days after deletion. What should you do?

- A. Run the Add-AzKeyVaultKey cmdlet.
- B. Run the `az keyvault update --enable-soft-delete true --enable-purge-protection true` CLI.
- C. Implement virtual network service endpoints for Azure Key Vault.
- D. Run the `az keyvault update --enable-soft-delete false` CLI.

**Suggested Answer: B**

When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.

Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.

When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.

The default retention period is 90 days, but it is possible to set the retention policy interval to a value from 7 to 90 days through the Azure portal. Once the retention policy interval is set and saved it cannot be changed for that vault.

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/overview-soft-delete>

Community vote distribution

B (100%)

 **anu\_ezio**  2 years, 8 months ago

**Selected Answer: B**

Correct Answer

upvoted 8 times

 **ca1502**  8 months ago

**Selected Answer: B**

<https://learn.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview>

upvoted 1 times

 **forgetfulalligator** 9 months, 1 week ago

**Selected Answer: B**

The `--enable-soft-delete` flag enables soft-delete on the key vault, which means that when objects are deleted, they are retained in a recoverable state for a specified period (90 days, in this case), rather than being immediately deleted.

The `--enable-purge-protection` flag turns on purge protection, which prevents the permanent deletion of objects until the soft-delete retention period has expired. This ensures that deleted objects can be recovered within the retention period.


upvoted 4 times

 **Ciupaz** 9 months, 1 week ago

**Selected Answer: B**

Soft-delete means that deleted key vault objects can be recovered for 90 days after deletion. Purge protection means that deleted key vault objects cannot be permanently deleted until the purge protection period has expired.

upvoted 1 times

 **harbox** 9 months, 1 week ago

**Selected Answer: B**

To configure an Azure Key Vault to allow recovery of deleted vaults or objects, you need to enable soft delete and purge protection.

Soft delete is a feature that allows you to recover a deleted Key Vault or its objects from the backup for a period of time, without the need to restore from a backup. By default, soft delete is disabled in Azure Key Vault.

Purge protection is a feature that protects Key Vault or its objects from being purged permanently, even if soft delete is enabled. This feature allows for recovery of the deleted object beyond the soft delete retention period. Purge protection is also disabled by default.

To enable soft delete and purge protection, you can use the Azure CLI command "az keyvault update" with the following options:

--enable-soft-delete true: This option enables soft delete for the Key Vault.

--enable-purge-protection true: This option enables purge protection for the Key Vault.

upvoted 1 times

🗳️ 👤 **viking2** 2 years ago

**Selected Answer: B**

B is correct. Soft-delete and purge protection are both needed, and B shows the correct syntax to achieve this

upvoted 1 times

🗳️ 👤 **uffuchsi** 2 years, 4 months ago

B. Purge defines how long a key vault entry is 'alive' before it can be deleted. Soft delete allows you to recover mistakenly deleted key vault entries.

upvoted 1 times

🗳️ 👤 **alexein74** 2 years, 5 months ago

B. Run the az keyvault update --enable-soft-delete true --enable-purge-protection true CLI.

To configure the application to allow recovery of an accidental deletion of the key vault or key vault objects for 90 days after deletion, you should run the az keyvault update --enable-soft-delete true --enable-purge-protection true CLI command. This command enables both soft delete and purge protection for the key vault.

Soft delete allows deleted key vaults and their contents to be retained and recoverable for the specified number of days after deletion.

Purge protection prevents key vault from being permanently deleted before the soft-delete retention period has elapsed.

upvoted 2 times

🗳️ 👤 **HafizSalmanMalik** 2 years, 6 months ago

Correct option is

Answer: B

upvoted 1 times

🗳️ 👤 **Topic9** 2 years, 6 months ago

Correct answer

upvoted 1 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

correct

upvoted 2 times

🗳️ 👤 **ahmad123456789** 2 years, 7 months ago

correct

upvoted 2 times

🗳️ 👤 **Ganeshd7** 2 years, 8 months ago

correct

upvoted 1 times

🗳️ 👤 **RachidBinAhmed** 2 years, 9 months ago

correct

upvoted 1 times

🗳️ 👤 **finnishr** 2 years, 9 months ago

Correct

upvoted 3 times

🗳️ 👤 **kampatra** 2 years, 9 months ago

Correct

upvoted 1 times

HOTSPOT -

You have developed a Web App for your company. The Web App provides services and must run in multiple regions.

You want to be notified whenever the Web App uses more than 85 percent of the available CPU cores over a 5 minute period. Your solution must minimize costs.

Which command should you use? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

az monitor metrics alert create -n myAlert -g myResourceGroup  
--scopes targetResourceID --condition " > 85"

	5m
--window size	CPU Usage
--evaluation-frequency	Percentage CPU
--auto-mitigate	avg Percentage CPU

### Answer Area

**Suggested Answer:**

az monitor metrics alert create -n myAlert -g myResourceGroup  
--scopes targetResourceID --condition "

CPU Usage  
Percentage CPU  
avg Percentage CPU

 > 85"

5m

--window size

--evaluation-frequency

--auto-mitigate

Reference:

<https://docs.microsoft.com/sv-se/cli/azure/monitor/metrics/alert>

 **2bruhornot2bruh** Highly Voted 9 months, 1 week ago

For anyone wondering why it's `--window-size` and not `--evaluation-frequency`: you want the average across 5 minutes. With `--evaluation-frequency` you don't go for averages, you simply check what the given value is at specific intervals.

upvoted 50 times

  **CellCS** Highly Voted  3 years, 2 months ago

Correct, Eng reference is here <https://docs.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest>

upvoted 20 times

  **Sandy4912** Most Recent 7 months, 1 week ago


Example Scenario:

Suppose you want to create a metric alert for when the CPU usage of your Web App exceeds 85% over a 5-minute window, and you want Azure to check this every 1 minute.

--window-size: 5 minutes (you want the CPU usage to exceed 85% for a full 5-minute period before triggering the alert).

--evaluation-frequency: 1 minute (Azure will check the CPU usage every minute to see if the 5-minute period still holds).

upvoted 2 times

 **zreaf29** 9 months, 1 week ago

Window size is the period of time over which Azure Monitor will collect metrics for a metric alert rule.

Evaluation frequency is the frequency with which Azure Monitor will evaluate the metric alert rule for violations.

if you set the window size to 5 minutes and the evaluation frequency to 1 minute, Azure Monitor will collect metrics for the previous 5 minutes and then evaluate the metric alert rule for violations every minute. This means that if there is a violation in the 5 minute window, Azure Monitor will detect



it within 1 minute.

We want to get percent "over a 5 minute period", we have to use window size parameter. Only use evaluation frequency parameter is not enough for getting question request.

upvoted 8 times

🗨️ **james2033** 1 year, 2 months ago

```
az monitor metrics alert create -n alert1 -g {ResourceGroup} --scopes {VirtualMachineID} --condition "avg Percentage CPU > 90" --window-size 5m --evaluation-frequency 1m --action "/subscriptions/<subscriptionId>/resourceGroups/<resourceGroupName>/providers/Microsoft.Insights/actionGroups/<actionGroupName>" apiKey={APIKey} type=HighCPU --description "High CPU"
```

<https://learn.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest#az-monitor-metrics-alert-create-examples>

upvoted 1 times

🗨️ **yunqing** 1 year, 9 months ago

Hello, may I know why Box 1 is 'avg Percentage CPU ', instead of 'Percentage CPU '. It wants to be notified whenever over than 85. Whenever, meaning the maximum, not the average?

upvoted 6 times

🗨️ **RDTAus** 1 year, 7 months ago

I am with you on this one, I think it should be Percentage CPU instead

<https://learn.microsoft.com/en-us/cli/azure/vm/monitor/metrics?view=azure-cli-latest>

upvoted 3 times

🗨️ **martin96** 1 year, 3 months ago

With the average over 5 minutes it would aggregate the cpu usage the last 5 min to provide an average and only alert if the average is over 85. Given that the question states "over a 5 min window" the Percentage CPU then if the CPU ever goes above 85 in a given 5 min window an alert would be sent. Say the CPU is around 50 and spikes for 1 tick to 86 then it would create an alert even though the average is 55.

upvoted 4 times

🗨️ **asdfdude** 1 year, 2 months ago

"uses more than 85 percent of the available CPU cores over a 5 minute period." The wording is so bad. Its could mean any of over 85 for the entire 5 minutes, over 85 at any point during the 5 minute windows, or avg > 85. My interpretation is cpu >85% for the entire duration, but there's no way to fulfill that requirement so the closest is average.

upvoted 1 times

🗨️ **uffuchsi** 2 years, 4 months ago

1. avg Percentage CPU, 2. window size

upvoted 1 times

🗨️ **Topic9** 2 years, 6 months ago

```
az monitor metrics alert create -n alert1 -g {ResourceGroup} --scopes {VirtualMachineID} \
--condition "avg Percentage CPU > 90" --window-size 5m --evaluation-frequency 1m \
--action "/subscriptions/<subscriptionId>/resourceGroups/<resourceGroupName>/providers/Microsoft.Insights/actionGroups/<actionGroupName>" \
apiKey={APIKey} type=HighCPU \
--description "High CPU"
```

upvoted 2 times

🗨️ **narenazure** 2 years, 6 months ago

Box 2: --window-size

Time over which to aggregate metrics in "##h##m##s" format.

<https://learn.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest>

upvoted 1 times

🗨️ **P000JAAAAAAAAA** 2 years, 6 months ago

Box 2: Window Size

```
az monitor metrics alert create -n alert1 -g {ResourceGroup} --scopes {VirtualMachineID} \
--condition "avg Percentage CPU > 90" --window-size 5m --evaluation-frequency 1m \
--action "/subscriptions/<subscriptionId>/resourceGroups/<resourceGroupName>/providers/Microsoft.Insights/actionGroups/<actionGroupName>" \
apiKey={APIKey} type=HighCPU \
--description "High CPU"
```

upvoted 1 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

what is the correct value for box2?

upvoted 1 times

🗨️ 👤 **DiveyK** 2 years, 9 months ago

--window-size : Time over which to aggregate metrics in "##h##m##s" format.

upvoted 3 times

🗨️ 👤 **RachidBinAhmed** 2 years, 9 months ago

the param --window size has a typo, missing a dash?

upvoted 1 times

🗨️ 👤 **srinz** 2 years, 9 months ago

Given answer is correct

<https://learn.microsoft.com/en-us/cli/azure/monitor/metrics/alert?view=azure-cli-latest#az-monitor-metrics-alert-create>

```
az monitor metrics alert create -n alert1 -g {ResourceGroup} --scopes {VirtualMachineID} \  
--condition "avg Percentage CPU > 90" --window-size 5m --evaluation-frequency 1m \  
--action "/subscriptions/<subscriptionId>/resourceGroups/<resourceGroupName>/providers/Microsoft.Insights/actionGroups/<actionGroupName>" \  
apiKey={APIKey} type=HighCPU \  
--description "High CPU"
```

upvoted 2 times

🗨️ 👤 **kampatra** 2 years, 9 months ago

Box 1: avg Percentage CPU > 85

Box 2: --evaluation-frequency 5m

upvoted 6 times

🗨️ 👤 **JohnSmith93** 2 years, 8 months ago

I would also think this is the solution because it reduces the times the query needs to be run, which is related to one to the criteria's in the question: "Your solution must minimize costs."

upvoted 1 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

why evaluation-frequency 5m, still not clear

upvoted 3 times

🗨️ 👤 **vadimprokopchuk** 2 years, 10 months ago

Box 1: avg Percentage CPU > 85

Box 2: --evaluation-frequency 5m

upvoted 3 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result. Establish if the solution satisfies the requirements.

You are configuring a web app that delivers streaming video to users. The application makes use of continuous integration and deployment.

You need to ensure that the application is highly available and that the users' streaming experience is constant. You also want to configure the application to store data in a geographic location that is nearest to the user.

Solution: You include the use of Azure Redis Cache in your design.

Does the solution meet the goal?

A. Yes


B. No

**Suggested Answer: B**

Community vote distribution

B (80%)

A (20%)

 **NStanhope** Highly Voted 3 years, 11 months ago

Correct - Instead, use CDN - <https://docs.microsoft.com/en-us/azure/architecture/best-practices/cdn>  
upvoted 39 times

 **StephanieMii** 3 years, 11 months ago

CDN might be best practice but the question is if it meets the goal. I think it does.  
upvoted 5 times

 **ranjitklive** 3 years, 11 months ago

Azure Cache for Redis is Distributed, in-memory, scalable solution providing super-fast data access. I think it is more suitable for database query results or session data caching and not content like video/images which you just fetch from azure storage and display.  
upvoted 12 times

 **Veks** 3 years, 3 months ago

Hmm, Azure Cache for Redis 6.0 includes Redis streams now...

I agree that CDN is correct because u have to ask your self who is writing these questions and why.

I'm sure that Redis is the solution that meets the goal, but the problem is, is it a solution form Microsoft exam question writer :)

upvoted 7 times

 **ucsdmiami2020** 3 years, 6 months ago

Further supporting the use of a CDN (Content Distributed Network), per Microsoft docs URL <https://docs.microsoft.com/en-us/azure/architecture/best-practices/cdn#how-and-why-a-cdn-is-used>

"Streaming video files to the client on demand. Video benefits from the low latency and reliable connectivity available from the globally located datacenters that offer CDN connections. Microsoft Azure Media Services (AMS) integrates with Azure CDN to deliver content directly to the CDN for further distribution."

upvoted 12 times

 **SagarUnhale** Most Recent 1 month, 2 weeks ago

**Selected Answer: B**

Answer : No

Azure Redis Cache is more focused on caching and improving performance rather than providing the necessary content delivery for streaming video, high availability, or geographic data storage.

upvoted 1 times

 **Fabiann** 9 months, 1 week ago

This is an open question :

I am agree about the use of Azure CDN about the videos management.

But the requirement is : "You need to ensure that the application is highly available".

There's nothing about cost etc, so why couldn't both Redis cache & CDN be included in the design ? CDN for videos management and Redis for other

stuff.

In any case, I will chose B(No) but I was curious to have advice.

upvoted 1 times

🗨️ **v\_arun** 1 year ago

**Selected Answer: A**

Reference:

<https://docs.microsoft.com/en-in/azure/cdn/>

upvoted 2 times

🗨️ **Santhosh\_Kumar** 1 year, 10 months ago

**Selected Answer: B**

Redis caches won't support streaming video content. Azure CDN is suitable.

upvoted 2 times

🗨️ **longnguyendh** 1 year, 11 months ago

The answer should be Azure Content Delivery Network (CDN)

upvoted 1 times

🗨️ **sayanchq** 2 years ago

Should be CDN .. Hence No

upvoted 1 times

🗨️ **UkadikeToch** 2 years ago

While Redis Cache can improve the performance and scalability of your web application, it may not directly address all the requirements of ensuring high availability and delivering a constant streaming experience to users. However, Redis Cache can be a valuable component within the overall architecture.

upvoted 2 times

🗨️ **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: B**

To store data in the geographic location nearest to the user, you can use Azure Content Delivery Network (CDN) to cache and deliver the video content. Azure CDN uses a network of edge servers that are distributed globally to cache and deliver content from the location nearest to the user.

upvoted 2 times

🗨️ **uffuchsi** 2 years, 4 months ago

B. No - Redis Stores / Caches backend data. Since this will need to Cache video content, use CDN

upvoted 1 times

🗨️ **gusches** 2 years, 5 months ago

**Selected Answer: B**

Microsoft CDN. URL <https://docs.microsoft.com/en-us/azure/architecture/best-practices/cdn#how-and-why-a-cdn-is-used>

upvoted 1 times

🗨️ **alexein74** 2 years, 5 months ago

B. No

Using Azure Redis Cache alone would not provide the level of high availability and constant streaming experience that you are looking for. Redis Cache is a in-memory data store that can improve the performance of web applications by caching data that is frequently accessed. However, it does not provide geographic replication or automatic failover to ensure high availability.

upvoted 2 times

🗨️ **garbas** 1 year, 8 months ago

Redis Cache has geographic replication:

<https://learn.microsoft.com/en-us/azure/azure-cache-for-redis/cache-how-to-active-geo-replication>

I think the key is that Redis is intended for database data, not for video streaming (too much data for in memory cache)

upvoted 1 times

🗨️ **Sukon\_Desknot** 2 years, 5 months ago

If you have the money, Azure Redis Cache combined with CDN is an excellent fit

upvoted 1 times

🗨️ **Topic9** 2 years, 6 months ago

CDN is right fit here

upvoted 1 times

  **Nedum1983** 2 years, 7 months ago

This is correct as it is stated clearly on the "how and why" section here: - <https://docs.microsoft.com/en-us/azure/architecture/best-practices/cdn>

Streaming video files to the client on demand. Video benefits from the low latency and reliable connectivity available from the globally located datacenters that offer CDN connections.


upvoted 1 times

  **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

use CDN

upvoted 1 times

  **\_philomath\_** 2 years, 8 months ago

CDN I think.

upvoted 1 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.

Establish if the solution satisfies the requirements.

You are configuring a web app that delivers streaming video to users. The application makes use of continuous integration and deployment.

You need to ensure that the application is highly available and that the users' streaming experience is constant. You also want to configure the application to store data in a geographic location that is nearest to the user.

Solution: You include the use of an Azure Content Delivery Network (CDN) in your design.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: A**

Reference:

<https://docs.microsoft.com/en-in/azure/cdn/>

Community vote distribution

A (100%)

 **alexsein74** Highly Voted 9 months, 1 week ago

**Selected Answer: A**

For a web app that delivers streaming video to users, you would need to use a combination of technologies that provide high availability, constant streaming experience and store data in a geographic location that is nearest to the user.

Azure Content Delivery Network (CDN) can be used to deliver streaming video content to users with low latency by caching the video at edge locations that are nearest to the user. Azure Media Services can be used to encode, package, and stream video content. Azure Media Services also provides built-in redundancy and failover options to ensure high availability.

Another option is to store the video files in Azure Blob storage, which provides a high-availability and low-latency storage solution that can be replicated to multiple regions.


upvoted 6 times

 **Ira89** Most Recent 3 months, 2 weeks ago

**Selected Answer: A**

This is right answer

upvoted 1 times

 **harbox** 1 year, 3 months ago

**Selected Answer: A**

Correct

upvoted 2 times

 **uffuchsi** 2 years, 4 months ago

A. Yes - Since this will need to Cache video content, use CDN

upvoted 2 times

 **guschess** 2 years, 5 months ago

**Selected Answer: A**

Check the former question

upvoted 1 times

 **murdokio** 2 years, 7 months ago

**Selected Answer: A**

agreed

upvoted 2 times

 **Nedum1983** 2 years, 7 months ago

Yes! I agreed.

upvoted 1 times

 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: A**

Correct

upvoted 1 times

🗲️ 👤 **Magis** 2 years, 7 months ago

**Selected Answer: A**

Correct.

upvoted 2 times

🗲️ 👤 **sfofana** 2 years, 8 months ago

CDN IS CORRECT!

upvoted 1 times

🗲️ 👤 **ranjitklive** 3 years, 11 months ago

What is the difference between Azure CDN and Azure Media Services?

<https://docs.microsoft.com/en-us/azure/media-services/latest/media-services-overview>

upvoted 3 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.

Establish if the solution satisfies the requirements.

You are configuring a web app that delivers streaming video to users. The application makes use of continuous integration and deployment.

You need to ensure that the application is highly available and that the users' streaming experience is constant. You also want to configure the application to store data in a geographic location that is nearest to the user.

Solution: You include the use of a Storage Area Network (SAN) in your design.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Community vote distribution

B (100%)

 **forgetfulalligator** Highly Voted 1 year, 10 months ago

**Selected Answer: B**


A Storage Area Network (SAN) is a dedicated network that provides access to consolidated, block-level data storage. It is used to increase the availability of data and improve the performance of applications that require access to shared data. However, SANs do not provide features to ensure high availability of web applications, nor do they provide a mechanism to store data in the geographic location nearest to the user.

upvoted 24 times

 **jvyas** Highly Voted 3 years, 2 months ago

Instead you use CDN


upvoted 12 times

 **harbox** Most Recent 9 months, 4 weeks ago

**Selected Answer: B**

The use of a Storage Area Network (SAN) alone is not sufficient to meet the stated requirements of ensuring high availability, improving the streaming experience, and storing data in a geographically nearest location. Other Azure services such as Azure Traffic Manager and Azure Media Services, as well as the selection of appropriate Azure regions, would need to be included in the overall solution design.


upvoted 2 times

 **deathRac3** 1 year, 7 months ago

**Selected Answer: B**

Its correct

upvoted 2 times

 **uffuchsi** 1 year, 10 months ago

B. No - SAN is a network architecture and won't solve the problem

upvoted 2 times

 **gus chess** 1 year, 11 months ago

**Selected Answer: B**

CDN instead

upvoted 1 times

 **alexein74** 1 year, 11 months ago

**Selected Answer: B**

B. No

A Storage Area Network (SAN) is a specialized, high-speed network that provides block-level access to data storage. SANs are mainly used to enhance the performance and availability of data storage. While a SAN can provide high-speed access to stored data, it would not meet the goal of ensuring high availability and constant streaming experience for the web app. SANs also do not provide geographical replication or automatic failover to ensure high availability.

upvoted 4 times

 **HafizSalmanMalik** 2 years ago



Correct answer is by using CDN

upvoted 1 times

🗨️ 👤 **Topic9** 2 years ago

CDN here

upvoted 1 times

🗨️ 👤 **BrainPowder** 2 years, 2 months ago

Ok CDN, makes sense

upvoted 2 times

🗨️ 👤 **Migsmiku** 2 years, 2 months ago

**Selected Answer: B**

Should use CDN

upvoted 1 times

🗨️ 👤 **upadhyayavi** 2 years, 8 months ago

What about ensuring the availability of the application??

upvoted 1 times

🗨️ 👤 **Veks** 2 years, 9 months ago

What storage u use has nothing to do with what your app is doing. OK, you will not use cold storage, of course, but faster or little slower disks don't have a big impact on the solution that has to provide a lot of content over internet and fast. This question is not about disk (storage) latency, but about internet access to content and its latency.

Definitely NO.

<https://www.backblaze.com/blog/whats-the-diff-nas-vs-san/>

upvoted 4 times

🗨️ 👤 **Vrushalij** 3 years, 4 months ago

No is correct answer

upvoted 4 times

🗨️ 👤 **ranjitklive** 3 years, 4 months ago

Does Azure storage uses SAN behind-the-scenes? The closes option in Azure for the requirement given is Azure Media Services.

<https://docs.microsoft.com/en-us/azure/media-services/latest/media-services-overview>

upvoted 2 times

🗨️ 👤 **RikinPatel** 1 year, 3 months ago

Azure Media Services will be retired June 30th, 2024.

upvoted 1 times

You develop a Web App on a tier D1 app service plan.

You notice that page load times increase during periods of peak traffic.

You want to implement automatic scaling when CPU load is above 80 percent. Your solution must minimize costs.

What should you do first?

- A. Enable autoscaling on the Web App.
- B. Switch to the Premium App Service tier plan.
- C. Switch to the Standard App Service tier plan.
- D. Switch to the Azure App Services consumption plan.

**Suggested Answer: C**


Configure the web app to the Standard App Service Tier. The Standard tier supports auto-scaling, and we should minimize the cost. We can then enable autoscaling on the web app, add a scale rule and add a Scale condition.

Reference:

<https://docs.microsoft.com/en-us/azure/monitoring-and-diagnostics/monitoring-autoscale-get-started> <https://azure.microsoft.com/en-us/pricing/details/app-service/plans/>


Community vote distribution

C (100%)

 **examTaker455** Highly Voted 3 years, 5 months ago

D1 is Shared Tier and does not offer autoscaling. You need to switch to Standard, Premium or Isolated.

upvoted 44 times

 **ZodiaC** 3 years, 5 months ago

So you have to reduce cost som you need standard app tier. 1000% correct!

upvoted 13 times

 **Gautam47** Highly Voted 3 years, 3 months ago

Tier D1 is a basically shared app service plan, so we need to move standard or premium plan to enable auto scaling. As we need to provide low cost solution, then standard plan will be best for this approach

upvoted 19 times


 **oskx2** Most Recent 9 months, 1 week ago

This answer may not be correct anymore? From microsoft pricing page: <https://azure.microsoft.com/en-us/pricing/details/app-service/linux-previous/>

The Standard service plan is still available for running less demanding production workloads. Pricing is based on the size and number of instances you run. For greater savings and higher performance, we strongly recommend using the latest App Service Premium plan.

Even still Standard is cheaper but you get less ram/storage. But standard pricing tier is being removed in favor of Premium.

upvoted 1 times

 **zreaf29** 1 year, 3 months ago

The D1 tier plan is a basic tier plan that does not offer autoscaling. For using autoscaling, you have to change D1 basic tier plan to Others(Standard, Premium, Isolated, Consumption). The Standard plan is cheapest plan which has auto scaling. So Answer is C.

upvoted 3 times

 **seanishok** 1 year, 4 months ago

Selected Answer: C

Standard Tier is the correct answer.

upvoted 1 times

 **longnguyendh** 1 year, 5 months ago

D1 is basic tier and It does not support auto scaling feature. => we should upgrade to Standard tier

upvoted 1 times

 **arpannaprap** 1 year, 5 months ago

Standard tier is the correct one.

upvoted 1 times

🗨️ **viking2** 1 year, 6 months ago

C is correct.

Auto scaling is supported on both Premium and Standard, but Standard is lower cost

upvoted 1 times

🗨️ **LuchoGonzalez** 1 year, 6 months ago

Enabling autoscaling (option A) can still be a valuable addition to further optimize the resource allocation based on the actual demand. By combining autoscaling with the upgraded Standard App Service tier, you can automatically scale the resources up or down based on CPU load, ensuring optimal performance during peak traffic while minimizing costs during periods of lower traffic.

Therefore, the recommended course of action in this scenario would be to switch to the Standard App Service tier (option C) and then enable autoscaling (option A) to effectively address the page load time issue during peak traffic while considering cost optimization.

upvoted 1 times

🗨️ **adilkhan** 1 year, 9 months ago

why not D Consumption plan? <https://learn.microsoft.com/en-us/azure/azure-functions/consumption-plan> The Consumption plan scales automatically, even during periods of high load. When running functions in a Consumption plan, you're charged for compute resources only when your functions are running. On a Consumption plan, a function execution times out after a configurable period of time.

upvoted 1 times

🗨️ **tom112** 1 year, 7 months ago

Because the consumption plan is for the Azure Function, not for a web app.

upvoted 4 times

🗨️ **forgetfulalligator** 1 year, 10 months ago

**Selected Answer: C**

the first step is to upgrade the App Service plan to the Standard or Premium tier, if it is not already on one of these tiers, and then enable autoscaling on the Web App

upvoted 1 times

🗨️ **wrak** 1 year, 10 months ago

Auto scale is NOT supported in tiers: Free (F1), Shared (D1 - only Windows) and Basic (B1, B2 and B3).

upvoted 3 times

🗨️ **uffuchsi** 1 year, 10 months ago

B. Switch to Standard as D1 does not allow autoscale. Standard is the cheapest option that supports autoscale.

upvoted 1 times

🗨️ **uffuchsi** 1 year, 10 months ago

Correction: C. Switch to Standard, not B

upvoted 1 times

🗨️ **guschess** 1 year, 11 months ago

**Selected Answer: C**

First at all, your choices are Standard or Premium. Reduce cost => Standard.

upvoted 1 times

🗨️ **santoshsidnal** 2 years ago

I agree that we need to enable auto-scaling to reduce the cost. But the question is "What you should do first?" hence switching to the standard app service makes more sense

upvoted 1 times

🗨️ **OPT\_001122** 2 years, 1 month ago

**Selected Answer: C**

C. Switch to the Standard App Service tier plan.

upvoted 1 times

🗨️ **OPT\_001122** 2 years, 1 month ago

**Selected Answer: C**

switch to Standard

upvoted 2 times

Your company's Azure subscription includes an Azure Log Analytics workspace.

Your company has a hundred on-premises servers that run either Windows Server 2012 R2 or Windows Server 2016, and is linked to the Azure Log Analytics workspace. The Azure Log Analytics workspace is set up to gather performance counters associated with security from these linked servers.

You must configure alerts based on the information gathered by the Azure Log Analytics workspace.

You have to make sure that alert rules allow for dimensions, and that alert creation time should be kept to a minimum. Furthermore, a single alert notification must be created when the alert is created and when the alert is resolved.

You need to make use of the necessary signal type when creating the alert rules.

Which of the following is the option you should use?

- A. The Activity log signal type.
- B. The Application Log signal type.
- C. The Metric signal type.
- D. The Audit Log signal type.

**Suggested Answer: C**

Metric alerts in Azure Monitor provide a way to get notified when one of your metrics cross a threshold. Metric alerts work on a range of multi-dimensional platform metrics, custom metrics, Application Insights standard and custom metrics.

Note: Signals are emitted by the target resource and can be of several types. Metric, Activity log, Application Insights, and Log.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-metric>

Community vote distribution

C (100%)

 **MiraA**  3 years, 2 months ago

Signal types:

- \* Activity Log ... includes service health records along with records on any configuration changes made to the resources (and is available to all Azure resources)
- \* Audit Log ... contains the history of sign-in activity and audit trail of changes made within a particular tenant
- \* Metric ... numerical values that are collected at regular intervals and describe some aspect of a system at a particular time
- \* Application Log ... ?

The Metrics feature can only store numeric data in a particular structure, whereas the Logs feature can store a variety of datatypes (each with its own structure).

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

<https://docs.microsoft.com/en-us/azure/azure-monitor/essentials/data-platform-metrics>

<https://docs.microsoft.com/en-us/azure/azure-monitor/logs/data-platform-logs>

upvoted 21 times

 **Pirgos**  3 years, 5 months ago

C, <https://docs.microsoft.com/en-us/azure/azure-monitor/alerts/alerts-metric-logs#configuring-metric-alert-for-logs>

upvoted 18 times

 **MiraA** 3 years, 2 months ago

From that page:

"Metric Alerts are stateful - only notifying once when alert is fired and once when alert is resolved; as opposed to Log alerts, which are stateless and keep firing at every interval if the alert condition is met."

upvoted 17 times

 **MiraA** 3 years, 2 months ago

From that page:

"Metric Alerts for Log provide multiple dimensions, allowing filtering to specific values like Computers, OS Type, etc. simpler; without the need for defining a complex query in Log Analytics."



upvoted 8 times

  **DheerajPatyal**  5 months ago

**Selected Answer: C**

Metric alerts are evaluated and triggered quickly because metrics are pre-aggregated and stored in Azure Monitor and also support notifications for both when an alert is triggered (fired) and when it is resolved (returned to normal)

upvoted 1 times

  **harbox** 9 months, 4 weeks ago

**Selected Answer: C**

Metric alerts in Azure Monitor provide a way to get notified when one of your metrics cross a threshold.

Metric alerts work on a range of multi-dimensional platform metrics, custom metrics, Application Insights standard and custom metrics.

Note: Signals are emitted by the target resource and can be of several types.

Metric, Activity log, Application Insights, and Log.

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/alerts-metric>



upvoted 1 times

  **forgetfulalligator** 1 year, 10 months ago

**Selected Answer: C**

The Metric signal type is the correct option for creating alert rules based on performance counters associated with security gathered by the Azure Log Analytics workspace. The Metric signal type allows you to create alerts based on metrics collected by Azure Monitor. This includes metrics collected by Log Analytics from your servers, such as performance counters associated with security. The Metric signal type allows you to set dimensions, which can be used to filter the metrics and reduce the noise of alerts. You can also configure the alert to be fired on multiple criteria and the alert creation time is kept to a minimum. Furthermore, a single alert notification is created when the alert is created and when the alert is resolved.

upvoted 3 times

  **wrak** 1 year, 10 months ago



There are two main requirements: (1) allow for dimensions, and (2) alert creation time should be kept to a minimum. Only Metric alert and Log alert supports split by dimensions, not Activity Log alerts. Additionally, we have to use KQL for Log alerts - this doesn't support 2nd requirement.

upvoted 2 times

  **uffuchsi** 1 year, 10 months ago

B. Metric

upvoted 1 times

  **uffuchsi** 1 year, 10 months ago

Correctio: C. Metric

upvoted 1 times


  **alexein74** 1 year, 11 months ago

**Selected Answer: C**

C. The Metric signal type

In order to configure alerts based on the information gathered by the Azure Log Analytics workspace, you should use the Metric signal type when creating the alert rules. This signal type allows you to create alerts based on the performance counters associated with security that are collected by the Azure Log Analytics workspace. Furthermore, Metric signals allow you to define alert rules using dimensions, which allows you to create more targeted alerts, and minimize the time spent creating alerts.

upvoted 2 times

  **nhannn** 2 years, 3 months ago

**Selected Answer: C**

C is correct. "Metric Alerts are stateful - only notifying once when alert is fired and once when alert is resolved; as opposed to Log alerts, which are stateless and keep firing at every interval if the alert condition is met."

upvoted 7 times

  **Anitay** 2 years, 3 months ago

C is correct answer

upvoted 1 times

You are developing a .NET Core MVC application that allows customers to research independent holiday accommodation providers. You want to implement Azure Search to allow the application to search the index by using various criteria to locate documents related to accommodation.

You want the application to allow customers to search the index by using regular expressions.

What should you do?

- A. Configure the SearchMode property of the SearchParameters class.
- B. Configure the QueryType property of the SearchParameters class.
- C. Configure the Facets property of the SearchParameters class.
- D. Configure the Filter property of the SearchParameters class.

**Suggested Answer: B**

The SearchParameters.QueryType Property gets or sets a value that specifies the syntax of the search query. The default is 'simple'. Use 'full' if your query uses the Lucene query syntax.

You can write queries against Azure Search based on the rich Lucene Query Parser syntax for specialized query forms: wildcard, fuzzy search, proximity search, regular expressions are a few examples.

Reference:


<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype>

Community vote distribution

B (100%)

 **7ack** Highly Voted 2 years, 12 months ago

Correct, although I'm not sure Azure search is in the exam...  
upvoted 22 times

 **TusharChandak** Highly Voted 11 months, 1 week ago

This is not part of AZ-204 anymore.  
upvoted 18 times

 **vizay** Most Recent 2 months, 1 week ago

**Selected Answer: B**

- SearchMode: Chooses between matching any or all keywords – not for regex.
  - QueryType: Enables advanced search features like regex when set to "full". ✓
  - Facets: Groups and counts results – used for filters like "by city."
  - Filter: Applies exact conditions (e.g., price < 100) – not for pattern matching.
  - ✓ Use QueryType to allow regex in Azure Search.
- upvoted 1 times

 **vizay** 2 months, 1 week ago

**Selected Answer: B**

SearchMode : This controls how results are matched:

"any" means match any word in the search

"all" means match all words

QueryType : This controls how smart or advanced your search can be.

You can set it to:

"simple" (default) – basic keyword search

"full" – allows advanced syntax like:

regex

wildcards

+, -, AND, OR, etc.

upvoted 1 times

 **N9** 9 months ago

**Selected Answer: B**

B is correct answer

upvoted 1 times

  **LuchoGonzalez** 1 year ago

B. Configure the QueryType property of the SearchParameters class.

The QueryType property of the SearchParameters class allows you to specify the type of query that will be executed against the Azure Search index. By setting the QueryType to "full" or "full+regex", you can enable regular expression support in the search queries.

upvoted 3 times

  **forgetfulalligator** 1 year, 4 months ago

**Selected Answer: B**



The QueryType property of the SearchParameters class specifies the type of query to use when searching the index. By default, it is set to Simple, which means that the query is interpreted as a simple query string. However, you can set it to Full, which means that the query is interpreted as a full Lucene query string, allowing for regular expressions to be used in the search query

upvoted 2 times

  **NombreFalso** 1 year, 4 months ago

huh, I'll ignore these

upvoted 1 times

  **wrak** 1 year, 4 months ago

out of the scope for AZ-204 exam; please remove/hide this question from the Question Bank

upvoted 6 times

  **uffuchsi** 1 year, 4 months ago

D. Filter property or B. QueryType property

upvoted 1 times

  **examinator007** 1 year, 5 months ago

Azure Search is not part of the AZ-204 - see: <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B>

upvoted 1 times

  **alexsein74** 1 year, 5 months ago



**Selected Answer: B**

B. Configure the QueryType property of the SearchParameters class.

To allow the application to search the index by using regular expressions, you should configure the QueryType property of the SearchParameters class in your .NET Core MVC application. The QueryType property allows you to specify the type of query to use when searching the index, including regular expression queries.

You can set the QueryType property to "full" to enable full-text search or "regex" to enable regular expression search.

upvoted 1 times

  **halfway** 1 year, 5 months ago

**Selected Answer: B**

QueryType

upvoted 1 times

  **OPT\_001122** 1 year, 10 months ago

<https://www.examttopics.com/exams/microsoft/az-204/view/6/>

same question #28 has answer D.

which one is correct

upvoted 1 times

  **Chang401** 1 year, 8 months ago

ask is different in both the questions.

upvoted 2 times

  **thedevguy** 1 year, 8 months ago



This one uses RegularExpressions while the other doesn't. Check again

upvoted 2 times

  **OPT\_001122** 1 year, 7 months ago

yes Thanks.



upvoted 1 times

  **gunz123** 2 years, 11 months ago

Answer is correct

<https://docs.microsoft.com/en-us/azure/search/search-query-overview>



upvoted 15 times

  **jay158** 2 years, 12 months ago

Answer : D

<https://docs.microsoft.com/en-us/azure/search/search-query-odata-filter#examples>

upvoted 1 times

  **jay158** 2 years, 11 months ago

my bad should be B

upvoted 7 times



You are a developer at your company.  
 You need to update the definitions for an existing Logic App.  
 What should you use?

- A. the Enterprise Integration Pack (EIP)
- B. the Logic App Code View
- C. the API Connections
- D. the Logic Apps Designer

**Suggested Answer: B**


Edit JSON - Azure portal -

1. Sign in to the Azure portal.
2. From the left menu, choose All services. In the search box, find "logic apps", and then from the results, select your logic app.
3. On your logic app's menu, under Development Tools, select Logic App Code View.
4. The Code View editor opens and shows your logic app definition in JSON format.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-overview> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions>

Community vote distribution



**abdou1987** Highly Voted 3 years, 11 months ago

the Logic App has not been part of the exam since March  
 upvoted 22 times

**MiraA** Highly Voted 3 years, 8 months ago

I think being a developer I can use both options - Logic App Code View or Logic Apps Designer. It depends on the change required, usage of a source control, the project complexity, CI/CD...

Note: Some Azure Logic Apps capabilities, such as defining parameters and multiple triggers in logic app definitions, are available only in JSON, not the Logic Apps Designer. So for these tasks, you must work in Code View or another editor.

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions>  
 upvoted 12 times

**rosireddy** Most Recent 2 months, 3 weeks ago

**Selected Answer: D**

The Logic Apps Designer is a visual tool within the Azure portal that allows you to easily view, edit, and update the definitions of an existing Logic App. It provides a user-friendly, drag-and-drop interface to modify workflows and integrate with various services.  
 upvoted 2 times

**DheerajPatyal** 5 months ago

**Selected Answer: B**

Logic App Code View in Logic Apps allows you to directly update the underlying JSON definition of a Logic App. This is the most precise method for updating definitions, especially if you need to make advanced changes or work directly with the Logic App schema.  
 upvoted 2 times

**harshadat** 11 months, 4 weeks ago

**Selected Answer: B**

Definitions are in the logic app code view  
 upvoted 2 times

**MaherEng** 1 year, 3 months ago

B. the Logic App Code View  
 upvoted 2 times

**Weam** 1 year, 6 months ago

**Selected Answer: B**

Logic app code view is specified for developers and enables them to get more control into making changes to an existing logic app however Logic app designer is mainly used by non developers which enables them to make changes without any coding changes

upvoted 2 times

🗨️ **ConanBarb** 1 year, 9 months ago

**Selected Answer: B**

That's an ignorant question. It totally depends on lots of things - not the least what type of change we are talking about. I would venture to state that it is virtually impossible, or at least takes at least 10x the time, to make a general big changes in code view such as adding new Actions, Connector Actions, a Loop or something. That's what you do in the designer with just a few button clicks and getting it right from start.

Small and specific changes yes, such as parameters, certain variable settings, etc

upvoted 3 times

🗨️ **UkadikeToch** 2 years ago

The logic app code view is in the logic apps designer. So i think the more accurate answer is THE LOGIC APP CODE VIEW

upvoted 2 times

🗨️ **moota** 2 years, 1 month ago

**Selected Answer: B**

ChatGPT: B: B. the Logic App Code View

Logic Apps in Azure provide a Code View that allows you to directly edit the JSON definition of your Logic App. This can be useful if you need to make advanced changes that are not easily done in the designer view, or if you are more comfortable working directly with the underlying code. The Logic Apps Designer is more for creating and configuring workflows in a visual manner, but the Code View is the more direct method to update the definitions.

upvoted 4 times

🗨️ **TheFivePips** 10 months, 2 weeks ago

funny that when I asked it, it says D

upvoted 1 times

🗨️ **NiebieskiEskimos** 2 years, 2 months ago

**Selected Answer: B**

Answer is correct

upvoted 1 times

🗨️ **adilkhan** 2 years, 3 months ago

D ChatGPT.

The Logic Apps Designer is a visual interface that allows you to design, configure, and manage Logic Apps. It provides a drag-and-drop interface that allows you to add and configure actions, triggers, and other components of your Logic App. You can use the Logic Apps Designer to modify the Logic App's workflow and its components, such as triggers, actions, and conditions.

upvoted 1 times

🗨️ **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: B**

B or D works for me, but I guess the keyword is "developer"?

upvoted 3 times

🗨️ **wrak** 2 years, 4 months ago

out-of-the-scope for the AZ-204 Exam; please remove from the Question Bank

upvoted 6 times

🗨️ **uffuchsi** 2 years, 4 months ago

D, or B for updating the Logic app, but since we need to update definitions, we can only use B

upvoted 1 times

🗨️ **Chang401** 2 years, 4 months ago

answer should be Logic App Code View.

To work with workflow definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, (go to the below link and just above the first note it's mentioned)

<https://learn.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions>

upvoted 1 times

🗨️ **Osip21** 2 years, 5 months ago

Selected Answer: D

D. the Logic Apps Designer.

upvoted 1 times

  **fkaracan** 2 years, 4 months ago

why? at least explain your answer...

upvoted 1 times

  **Molnmannen1337** 1 year, 11 months ago

The designer is what you usually work in if you develop a logic app. The code view is just a bunch of json code.

upvoted 2 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.

Establish if the solution satisfies the requirements.

You are developing a solution for a public facing API.

The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.

You must configure back-end authentication for the API Management service instance.

Solution: You configure Basic gateway credentials for the Azure resource.

Does the solution meet the goal?

A. Yes

B. No


**Suggested Answer: B**

API Management allows to secure access to the back-end service of an API using client certificates.

Reference:

<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity>

*Community vote distribution*



**lexowe9241** 3 years, 9 months ago

Hmm, there is "Authenticate with Basic policy" for api-management <https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#Basic>

So maybe Yes is correct answer

upvoted 22 times

**MiraA** 3 years, 8 months ago

The answer is YES (Basic + the HTTP(s) endpoint) based on this screenshot:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

Target: "Azure Logic App" or "HTTP(s) endpoint"

Gateway credentials: "None" or "Basic" or "Client cert"

upvoted 3 times

**MiraA** 3 years, 8 months ago

My own correction... :-(

The answer is NO.

The solution mentions "You configure Basic gateway credentials for the Azure resource." and there is no such option for generic Azure resource.

upvoted 11 times

**koolexam** 3 years, 6 months ago

But you can do that at collection level in APIM.

upvoted 1 times

**warchoon** 2 years, 4 months ago

Another tricky question about settings UI, that will soon be deprecated :)

upvoted 1 times

**basquiat** 1 year, 10 months ago

Authentication policy sounds quite different from authentication gateway IMHO.

upvoted 2 times

**PhilLI** 3 years, 5 months ago

Agree. My research lead me also to this page.

Authentication policies:

- Authenticate with Basic -> Authenticate with a backend service using Basic authentication.
  - Authenticate with client certificate -> Authenticate with a backend service using client certificates.
  - Authenticate with managed identity -> Authenticate with the managed identity for the API Management service.
- upvoted 8 times

  **Mev4953** 3 years, 5 months ago

I agree with you.

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>

upvoted 1 times

  **leonidn** Highly Voted 3 years, 5 months ago

Selected Answer: B

Basic client credentials can be defined for HTTP endpoint, not an App Service. My previous answer is incorrect.

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#Basic>

upvoted 16 times

  **Sassy1** Most Recent 2 months, 1 week ago

Selected Answer: B

B- No

Why this doesn't meet the goal:

- Azure API Management supports several more secure and appropriate back-end authentication mechanisms, such as:
- Managed identity (preferred for Azure App Service back ends)
- Client certificates
- OAuth 2.0

- Basic auth is not secure unless used over HTTPS, and is not recommended for authenticating with Azure App Services unless absolutely necessary and well-protected.

- It does not integrate directly with Azure Active Directory (AAD) or securely manage credentials like Managed Identity does.

upvoted 2 times

  **Surnok** 5 months, 3 weeks ago

Selected Answer: B

Why this does not meet the goal



Basic gateway credentials are primarily used for simple scenarios where the back-end service expects a username and password for authentication. For public-facing APIs, you typically need stronger and more secure authentication mechanisms, such as:

OAuth 2.0

JWT (JSON Web Tokens)

Managed Identity for Azure resources


upvoted 2 times

  **overhill** 10 months, 2 weeks ago

Selected Answer: A

they're asking for basic auth with user and password. APIM supports that, whether or not it is secure is another story... this is a yes

upvoted 1 times

  **MaherEng** 1 year, 3 months ago

Selected Answer: B

good answer is B

upvoted 1 times

  **[Removed]** 1 year, 5 months ago

Configuring Basic gateway credentials is not recommended for securing a public-facing API because it involves sending credentials (username and password) with each request. This method can expose security vulnerabilities, as credentials are transmitted in an easily readable format. It's better to use more secure methods like API keys, OAuth tokens, or other authentication mechanisms for better protection in a public API scenario

upvoted 6 times

  **Mick1z8** 1 year, 8 months ago

Selected Answer: B

There are two components we're interested in:

- APIM, and
- App Service (resource)

The App Service doesn't support basic-auth at all; though APIM does. The tricky part is the word "resource" which is App Service.

upvoted 1 times

🗨️ 👤 **Mick1z8** 1 year, 8 months ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/azure/api-management/authentication-basic-policy>

It clearly states (in the first line), that `authentication-basic` policy can be used to authenticate with a backend service.

upvoted 1 times

🗨️ 👤 **Vladimir\_Gajinov** 1 year, 9 months ago

**Selected Answer: A**

The correct answer is YES.

Use the authentication-basic policy to authenticate with a backend service using Basic authentication.

<https://learn.microsoft.com/en-us/azure/api-management/authentication-basic-policy>

upvoted 3 times

🗨️ 👤 **JH81** 2 years ago

**Selected Answer: B**

Yes you can use Basic but it certainly would not be a recommended solution because it's not secure for a public facing API. So using MS think the answer is No.

upvoted 1 times

🗨️ 👤 **BikashKatwal** 2 years ago

B. No

The solution does not meet the goal of configuring back-end authentication for the API Management service instance. Configuring Basic gateway credentials for the Azure resource does not provide authentication for the API Management service. Basic gateway credentials are used for authentication between the client and the API Management gateway, but it does not provide authentication for the back-end service hosted in the Azure App Service instance.

To configure back-end authentication for the API Management service instance, you would typically use authentication mechanisms such as OAuth, API keys, or client certificates to authenticate and secure the communication between the API Management gateway and the back-end service.

upvoted 4 times

🗨️ 👤 **Mattt** 2 years, 3 months ago

B. No.

Configuring Basic gateway credentials for the Azure resource would provide authentication for accessing the Azure resource itself, but it does not provide authentication for the API Management service instance.

To configure back-end authentication for the API Management service instance, you should use one of the following authentication options:

Client Certificate authentication

Token-based authentication

OAuth 2.0 authentication

These options provide secure authentication and access control for the API Management service instance and its associated APIs.

upvoted 6 times

🗨️ 👤 **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: A**

While Basic authentication is a supported option for API Management gateway authentication, it is not recommended for public-facing APIs due to its inherent security weaknesses. Basic authentication transmits user credentials in plain text, which can be easily intercepted and compromised.

upvoted 3 times

🗨️ 👤 **forgetfulalligator** 2 years, 4 months ago

Opps sorry meant to select B

upvoted 1 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

B. Would you not instead use Anon and authenticate pass through JWT token to the backend?

upvoted 2 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

So it will work much faster than certs.

upvoted 1 times

🗨️ 👤 **alexein74** 2 years, 5 months ago

**Selected Answer: B**

B. No

Configuring Basic gateway credentials for the Azure resource would not provide a secure solution for back-end authentication for the public facing API. Basic authentication sends the user's credentials in plain text over the network, making it vulnerable to eavesdropping and man-in-the-middle attacks. This is not suitable for public facing APIs.

A more secure solution would be to use OAuth2 or OpenID Connect protocol for back-end authentication. This allows the API Management service to securely authenticate the client against an identity provider, such as Azure Active Directory, and obtain an access token for the back-end service.

In summary, configuring Basic gateway credentials for the Azure resource does not meet the goal of securing back-end authentication for the public facing API.

upvoted 2 times

🗨️ 👤 **fkaran** 2 years, 4 months ago

not everytime chatgpt generates the correct answer. you need to do your own research before paste here...

upvoted 3 times

🗨️ 👤 **mitu7** 2 years, 6 months ago

**Selected Answer: B**

ChatGPT:

No, the solution of configuring Basic gateway credentials for the Azure resource does not meet the goal of configuring back-end authentication for the API Management service instance.

To configure back-end authentication for the API Management service instance, you need to specify the credentials that the API Management service instance should use to authenticate to the API back end. There are several options for back-end authentication, such as using an Azure Active Directory tenant, a certificate, or an API key. The Basic gateway credentials are used to authenticate the API Management service instance to the Azure resource, which is not the same as back-end authentication for the API Management service instance.

To properly configure back-end authentication for the API Management service instance, you need to choose the appropriate authentication method and specify the credentials accordingly. You can find more information about configuring back-end authentication in the Azure API Management documentation.

upvoted 2 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.

Establish if the solution satisfies the requirements.

You are developing a solution for a public facing API.

The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.

You must configure back-end authentication for the API Management service instance.

Solution: You configure Client cert gateway credentials for the HTTP(s) endpoint.

Does the solution meet the goal?

A. Yes

B. No


**Suggested Answer: B**

The API back end is hosted in an Azure App Service instance. It is an Azure resource and not an HTTP(s) endpoint.

Reference:

<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity>

*Community vote distribution*



🗳️ 👤 **jay158** Highly Voted 3 years, 12 months ago

This is scenario questions.

If backend is accepts HTTP(S)

Then Basic AUTH or Certificate will work.

so Client Certificate + HTTP(s) YES

upvoted 43 times

🗳️ 👤 **r3verse** 2 years, 7 months ago

APIM supports both, but that is only part of the question.

Does App Service itself support basic auth? NO not really (only a hacky way).

Does app Service Support certificate auth? YES, built in the Azure portal directly as a setting for app Service.

So my answer would be YES for client certificate.

upvoted 3 times

🗳️ 👤 **aradice** Highly Voted 3 years, 12 months ago

yes? <https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates>

upvoted 14 times

🗳️ 👤 **aradice** 3 years, 12 months ago

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#ClientCertificate>

upvoted 5 times

🗳️ 👤 **warchoon** 2 years, 4 months ago

It's the only link that alive. All others are dead and redirect to another pages or target deprecated sources.

upvoted 1 times

🗳️ 👤 **MiraA** 3 years, 8 months ago

The answer is YES (client cert + the HTTP(s) endpoint) based on this screenshot:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

Target: "Azure Logic App" or "HTTP(s) endpoint"

Gateway credentials: "None" or "Basic" or "Client cert"

upvoted 7 times

🗳️ 👤 **TheFivePips** Most Recent 10 months, 2 weeks ago



**Selected Answer: B**

The API back end is hosted in an Azure App Service instance. It is an Azure resource and not an HTTP(s) endpoint.

upvoted 1 times

  **Masombutronics** 1 year ago

The solution you've described does not meet the goal. While configuring client certificate gateway credentials for the HTTP(s) endpoint is a valid approach for securing the API, it doesn't directly address the requirement to configure back-end authentication for the API Management service instance.

To meet the goal, consider using Azure API Management policies to enforce authentication and authorization. Specifically, you can use policies like authentication-certificate or authentication-oauth2 to secure the API endpoints. These policies allow you to validate client certificates or use OAuth2 tokens for authentication.

Remember to configure the API Management service to validate incoming requests against the back-end API's certificate or OAuth2 token. This way, you'll ensure that only authorized clients can access your API.

upvoted 2 times

  **Shreyas2811** 1 year, 2 months ago

Yes, the solution meets the goal. By configuring Client cert gateway credentials for the HTTP(s) endpoint, you are implementing back-end authentication for the API Management service instance. This means that the API Management service will only allow access to the back-end hosted in the Azure App Service instance if the appropriate client certificate is presented. This provides a secure means of authenticating and authorizing requests from the API Management service to the back-end API.


upvoted 2 times

  **VK1989** 1 year, 2 months ago

**Selected Answer: B**

OAuth 2.0 or API Key is needed to secure public facing API as per API Management Service. This was also the answer from COPILOT so answer is B

upvoted 2 times

  **[Removed]** 1 year, 5 months ago

Yes. Configuring Client Certificate gateway credentials for the HTTP(s) endpoint is a more secure solution for back-end authentication in a public-facing API scenario. It involves using client certificates to authenticate the requests made to the API Management service instance, enhancing security compared to basic credentials.

upvoted 1 times

  **Weam** 1 year, 6 months ago

**Selected Answer: A**

Authentication policies

Authenticate with Basic - Authenticate with a backend service using Basic authentication.

Authenticate with client certificate - Authenticate with a backend service using client certificates.

Authenticate with managed identity - Authenticate with a backend service using a managed identity.

upvoted 1 times

  **BikashKatwal** 2 years ago

The solution mentioned, which is configuring Client cert gateway credentials for the HTTP(s) endpoint, does not meet the goal of configuring back-end authentication for the API Management service instance. Client certificate authentication is used for authenticating clients accessing the API Management service, not for authenticating the back-end service itself.

To configure back-end authentication for the API Management service instance hosted in an Azure App Service, you have several options, such as:

API key authentication: Generate an API key and configure it in the API Management service to authenticate requests sent to the back-end service.

OAuth 2.0 authentication: Configure OAuth 2.0 authentication between the API Management service and the back-end service, allowing clients to authenticate using OAuth 2.0 tokens.

Certificate-based authentication: If your back-end service supports certificate-based authentication, you can configure the API Management service to authenticate requests using client certificates.

Therefore, the correct answer is:

B. No

upvoted 7 times

🗨️ **ami2023** 1 year, 7 months ago

quoted from your reply "Certificate-based authentication: If your back-end service supports certificate-based authentication, you can configure the API Management service to authenticate requests using client certificates.", and this is listed as a way to configure back-end authentication.

upvoted 1 times

🗨️ **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates>

upvoted 1 times

🗨️ **uffuchsi** 2 years, 4 months ago

A. Yes

upvoted 1 times

🗨️ **AAcademy** 2 years, 4 months ago

I would say "YES". Secure backend services using client certificate authentication in Azure API Management: <https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates>

upvoted 1 times

🗨️ **alexein74** 2 years, 5 months ago

**Selected Answer: B**

B. No

Configuring client cert gateway credentials for the HTTP(s) endpoint would not provide a suitable solution for back-end authentication for the public facing API. Client certificate authentication is based on the possession of a certificate and private key, which can be installed on the client device. However, this approach is not suitable for public facing APIs, as it requires the client to have a unique certificate, which is not easy to manage, and it's not a scalable solution for public facing APIs.

upvoted 4 times

🗨️ **shacent** 2 years, 5 months ago

chatGPT: This solution does not meet the goal of configuring back-end authentication for the API Management service instance. Client certificate gateway credentials are used for client-side authentication, which is not the same as back-end authentication. Back-end authentication is used to authenticate the API Management service instance with the back-end service hosted in the Azure App Service instance.

For back-end authentication, you can use Azure Active Directory (AAD) authentication, Azure AD B2C, or OAuth 2.0 authentication to authenticate the API Management service instance with the back-end service.

upvoted 4 times

🗨️ **deepbh** 2 years, 6 months ago

A (Yes) is correct.

<https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates>

"API Management allows you to secure access to the backend service of an API using client certificates." ... "You should have your backend service configured for client certificate authentication. To configure certificate authentication in the Azure App Service, refer to this article."

Then see: <https://learn.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth?tabs=azurecli#enable-client-certificates>

upvoted 1 times

🗨️ **[Removed]** 2 years, 9 months ago

**Selected Answer: A**

Client certificate is supported. Even if I select an App service while creating the API it shows the back-end authentication option as Https endpoint and not as Azure resource.

upvoted 1 times

🗨️ **SaintBahamut** 3 years, 2 months ago

**Selected Answer: A**

Answer is Yes, its obvious if you check it yourself

Design Tab->Backend panel->HTTP(s) endpoint->And you have None, Basic and Certificate

upvoted 3 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.

Establish if the solution satisfies the requirements.

You are developing a solution for a public facing API.

The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.

You must configure back-end authentication for the API Management service instance.

Solution: You configure Basic gateway credentials for the HTTP(s) endpoint.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**


API Management allows to secure access to the back-end service of an API using client certificates. Furthermore, the API back end is hosted in an Azure App

Service instance. It is an Azure resource and not an HTTP(s) endpoint.

Reference:

<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity>

Community vote distribution



**jay158** Highly Voted 3 years, 12 months ago

This is scenario questions.

If backend is accepts HTTP(S)

Then Basic AUTH or Certificate will work.

so Basic + HTTPS Yes

upvoted 42 times

**MiraA** 3 years, 8 months ago

Target: "Azure Logic App" or "HTTP(s) endpoint"

Gateway credentials: "None" or "Basic" or "Client cert"

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

upvoted 7 times

**ndchris2003** Highly Voted 3 years, 11 months ago

The correct answer is Yes.

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#ClientCertificate>

upvoted 11 times

**Arijeet2712** Most Recent 6 days, 23 hours ago

Selected Answer: A

As the Basic Creds are supported with https endpoint

upvoted 1 times

**john\_doe\_951627384** 1 month, 3 weeks ago

Selected Answer: A

check the screenshot:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

upvoted 1 times

**TheFivePips** 10 months, 2 weeks ago

Selected Answer: B

Client Cert Gateway Credentials:

Can be used with both HTTP endpoints and Azure resources if both are configured to handle and validate client certificates.

Provides strong security and is suitable for high-security scenarios, but involves a more complex setup.

Basic Gateway Credentials:

Suitable for HTTP endpoints that support Basic Authentication but not suitable for Azure resources like Azure App Service.

Basic Authentication is less secure and generally not used for Azure resources, which typically require Azure AD, OAuth tokens, or managed identities for authentication.

For a public-facing API, client certificate authentication is generally preferred if both APIM and the backend API are properly configured to use and validate client certificates. Basic Authentication is simpler but less secure and not suitable for most Azure resources. Azure App Service does not support Basic Authentication (Basic gateway credentials) for securing access to the service itself. For securing Azure App Service, you should use Azure AD, OAuth 2.0, or other more secure authentication methods

upvoted 2 times

  **Shreyas2811** 1 year, 2 months ago

No, the solution does not meet the goal. Configuring Basic gateway credentials for the HTTP(s) endpoint provides authentication for clients accessing the API through the API Management gateway, but it does not specifically address back-end authentication for the API Management service instance. Back-end authentication typically involves authenticating requests between the API Management service instance and the back-end hosted in Azure App Service. Basic gateway credentials are more commonly used for authenticating clients accessing the API through the API Management gateway, rather than for securing communication between the gateway and the back-end service.

upvoted 1 times

  **rumbur4k** 1 year, 1 month ago

I think that you mix up client auth -> APIM and APIM -> Backend

<https://learn.microsoft.com/en-us/azure/api-management/authentication-basic-policy>

I would say A is correct.



upvoted 1 times

  **VK1989** 1 year, 2 months ago

**Selected Answer: B**

Answer is B - NO

upvoted 1 times

  **oskx2** 1 year, 4 months ago

**Selected Answer: A**

Answer is YES. Based on this: <https://learn.microsoft.com/en-us/azure/api-management/authentication-basic-policy>

upvoted 1 times

  **raymond\_abcd** 1 year, 4 months ago

**Selected Answer: B**

This is not the preferred way, Basic authentication is unsecure because of sending the username/password in the header. You must use TLS to protect the credentials. The preferred way is using a client certificate. Which can be re-used for all backends that are residing in the same api management service

upvoted 2 times

  **[Removed]** 1 year, 5 months ago



No. Configuring Basic gateway credentials for the HTTP(s) endpoint in a public-facing API scenario is not ideal for security reasons. Basic credentials involve sending a username and password with each request, which can expose sensitive information and is less secure. It's recommended to use more robust authentication mechanisms like API keys, OAuth tokens, or client certificates for enhanced security in such scenarios.

upvoted 1 times

  **manopeydakon** 1 year, 5 months ago

Correct answer is Yes, Admin change the answer: Configuring Basic gateway credentials for the HTTP(s) endpoint is a valid solution for back-end authentication in API Management. Basic authentication involves sending a username and password with each request.

upvoted 1 times

  **garbas** 1 year, 8 months ago

**Selected Answer: A**

API Management gateway supports basic authentication.

You have a RESTful service, so there is an HTTP(s) endpoint.

App Service doesn't have built-in support for basic authentication, but you can use the security features of the framework the service is built on (<https://learn.microsoft.com/en-us/azure/app-service/overview-authentication-authorization#why-use-the-built-in-authentication>)

The question only says "you must configure authentication". It doesn't say it has to be the best solution or the most secure. Maybe Microsoft is expecting No in this question, but the solution meets the goal, so the answer should be Yes.

upvoted 3 times

🗨️ 👤 **juanHuizar** 1 year, 9 months ago

No, configuring Basic gateway credentials for the HTTP(s) endpoint does not fully meet the goal of configuring backend authentication for the API Management service instance, especially if this API is public-facing and needs to be secure.

Basic gateway credentials refer to HTTP Basic Authentication, where a client sends a username and password with each request.

upvoted 3 times

🗨️ 👤 **TheMCT** 1 year, 9 months ago

**Selected Answer: B**

Configuring Basic gateway credentials for the HTTP(s) endpoint in API Management is not typically used for back-end authentication. Basic gateway credentials are used to secure access to the API Management gateway itself, allowing clients to authenticate when making requests to the API Management layer. It does not secure the communication between API Management and the back-end Azure App Service.

upvoted 1 times

🗨️ 👤 **BikashKatwal** 2 years ago

The solution mentioned, which is configuring Basic gateway credentials for the HTTP(s) endpoint, does meet the goal of configuring back-end authentication for the API Management service instance.

Configuring Basic gateway credentials means that the API Management service will authenticate requests sent to the back-end service using basic authentication. Basic authentication requires clients to include a username and password in the request headers.

By configuring Basic gateway credentials for the HTTP(s) endpoint, you can enforce authentication for incoming requests and ensure that only authorized clients can access the API back end hosted in the Azure App Service instance.

Therefore, the correct answer is:

A. Yes

upvoted 2 times

🗨️ 👤 **Net\_IT** 2 years, 4 months ago

**Selected Answer: A**

Answer should be YES according to this link: <https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

upvoted 1 times

🗨️ 👤 **Michael2023** 2 years, 5 months ago

Selected Answer: B

Because the Backend is an App Service the Microsoft security recommendations to protect back-end resources are: User/Application Identity and/or require client certificate authentication

APIM would support Basic & Client, but it also depends on which kind of backend application you are using

<https://learn.microsoft.com/en-us/azure/app-service/security-recommendations#identity-and-access-management>

upvoted 1 times

Note: The question is included in a number of questions that depicts the identical set-up. However, every question has a distinctive result.

Establish if the solution satisfies the requirements.

You are developing a solution for a public facing API.

The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end.

You must configure back-end authentication for the API Management service instance.

Solution: You configure Client cert gateway credentials for the Azure resource.

Does the solution meet the goal?

A. Yes

B. No


#### Suggested Answer: A

API Management allows to secure access to the back-end service of an API using client certificates.

Reference:

<https://docs.microsoft.com/en-us/rest/api/apimanagement/apimanagementrest/azure-api-management-rest-api-backend-entity>

Community vote distribution



**jay158** Highly Voted 3 years, 12 months ago

This is scenario questions.

If backend is accepts HTTP(S)

Then Basic AUTH or Certificate will work.

so Certificate + Azure Resource NO

<https://www.youtube.com/watch?v=HQ0U7lwP93o>

upvoted 33 times

**MiraA** 3 years, 8 months ago

Target: "Azure Logic App" or "HTTP(s) endpoint"

Gateway credentials: "None" or "Basic" or "Client cert"

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

upvoted 8 times

**PhilLI** 3 years, 5 months ago

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth>

So I think it's YES

upvoted 3 times

**warchoon** 2 years, 4 months ago

No. Your page is not about backend. It's about CLIENT certificate authentication

upvoted 1 times

**warchoon** 2 years, 4 months ago

Sorry my joke. But it is true.

upvoted 2 times

**Ciupaz** Highly Voted 1 year, 7 months ago

Hope don't find this question in my exam.

upvoted 12 times

**ThierryAhn** 1 year, 3 months ago

I hope too. Do you pass your exam ?

upvoted 1 times

**manvith7** Most Recent 5 months, 2 weeks ago

**Selected Answer: A**

You are developing a solution for a public facing API. The API back end is hosted in an Azure App Service instance. You have implemented a RESTful service for the API back end. You must configure back-end authentication for the API Management service instance. Solution: You configure Client cert gateway credentials for the Azure resource. Does the solution meet the goal?

Yes, the solution meets the goal. Configuring Client certificate gateway credentials for the Azure resource ensures secure back-end authentication for the API Management service instance. This method provides a robust way to authenticate and secure communication between the API Management gateway and the back-end service.

upvoted 1 times

🗳️ 👤 **TheFivePips** 10 months, 2 weeks ago

**Selected Answer: A**

This works because the backend is not an HTTP endpoint, its an Azure resource. It also works because App Service doesn't support basic auth out of the box, and it is considerably less secure even if it did.

upvoted 1 times

🗳️ 👤 **8ac3742** 11 months ago

The answer should be No, the client certificate is typically used for mutual authentication between the client application and the APIM gateway after APIM uploads the client certificate.

upvoted 1 times

🗳️ 👤 **Frawsty** 11 months, 1 week ago

Pasted this question into MS copilot and got this response

The solution of configuring Client Certificate Gateway credentials for the HTTP(s) endpoint does not meet the goal. While client certificate authentication provides an additional layer of security, it is not the recommended approach for authenticating the API Management service instance with the back-end hosted in Azure App Service.

upvoted 1 times

🗳️ 👤 **raymond\_abcd** 1 year, 4 months ago

**Selected Answer: A**

This is described on this page: <https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 5 months ago

No, the solution does not meet the goal. Configuring Client cert gateway credentials for the Azure resource is not a standard practice for back-end authentication in API Management. Typically, back-end authentication is configured at the API Management level, often involving the use of API keys, OAuth, or other authentication mechanisms specific to the API Management service.

upvoted 1 times

🗳️ 👤 **TheMCT** 1 year, 9 months ago

**Selected Answer: B**

The answer is; NO

Configuring "Client cert gateway credentials" for the Azure resource suggests that you are configuring client certificate authentication for the API Management gateway itself when accessing Azure resources, but it may not necessarily secure the communication between API Management and the back-end Azure App Service.

upvoted 4 times

🗳️ 👤 **raymond\_abcd** 1 year, 4 months ago

<https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates#configure-an-api-to-use-client-certificate-for-gateway-authentication>

upvoted 1 times

🗳️ 👤 **JH81** 2 years ago

**Selected Answer: A**

One of the questions is for securing with Basic and the other with a client cert. Clearly MS is saying Basic = bad and client cert = good.

upvoted 3 times

🗳️ 👤 **BikashKatwal** 2 years ago

A. Yes



The solution meets the goal of configuring back-end authentication for the API Management service instance. By configuring Client cert gateway credentials for the Azure resource, you can authenticate and secure the communication between the API Management gateway and the back-end service hosted in the Azure App Service instance.



Client certificate authentication involves using X.509 client certificates to verify the identity of the client requesting access to the API Management gateway. This authentication method ensures that only clients with valid certificates can access the back-end service.

By configuring Client cert gateway credentials, you establish a secure and authenticated connection between the API Management gateway and the Azure App Service instance, thereby meeting the requirement for back-end authentication.

upvoted 5 times

  **basquiat** 1 year, 10 months ago

If I got it right the link below shows how to achieve that. <https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates>

upvoted 2 times

  **deathRac3** 2 years, 1 month ago

**Selected Answer: A**

Looks right

upvoted 2 times



  **forgetfulalligator** 2 years, 4 months ago

**Selected Answer: A**

When client certificate authentication is used for the API Management gateway, a client must provide a client certificate in the header of each request to access the API. The client certificate is then validated by the gateway before forwarding the request to the back-end API. This can provide a high level of authentication and help prevent unauthorized access to the back-end API.



Client certificate authentication is a more secure mechanism than Basic authentication, as it uses a secure communication channel and cryptographic keys to authenticate clients. Additionally, it provides better scalability, as it does not require the gateway to maintain a large number of user credentials.

upvoted 4 times

  **mabdo** 2 years, 4 months ago

Are the website answers the official ones? This is getting confusing!!!

upvoted 4 times

  **joshfry** 2 years, 4 months ago

No. They don't have access to official answers. They try their best, but always read the discussion.

upvoted 3 times

  **Michael2023** 2 years, 5 months ago

Selected Answer: A -> yes

as long as a certificate is used to can authenticate to the backend via "Custom URL" or "Azure resource" (Function App, Logic App, Web App)... in both cases you provide an URL to which the requests are forwarded to.

From the App Service side you should use HTTPS only or User/System Managed Identity to authenticate to the App Service.

So that means, the only solutions for this scenario would be using a certificate for HTTP(s) endpoint and resources.

<https://learn.microsoft.com/en-us/azure/app-service/security-recommendations#identity-and-access-management>

<https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates>

upvoted 2 times

  **warchoon** 2 years, 4 months ago

Logic App is ok. But it is not the App Service instance. Others use gateway Https endpoint settings.

upvoted 1 times



  **alexein74** 2 years, 5 months ago

**Selected Answer: B**

B. No

Configuring client cert gateway credentials for the Azure resource would not provide a suitable solution for back-end authentication for the public facing API. Client certificate authentication is based on the possession of a certificate and private key, which can be installed on the client device. However, this approach is not suitable for public facing APIs, as it requires the client to have a unique certificate, which is not easy to manage, and it's not a scalable solution for public facing APIs.

upvoted 3 times

  **deepbh** 2 years, 6 months ago



**Selected Answer: B**

Azure Logic App and HTTP(S) Endpoint are the two choices on the portal page. Whenever the question is about using app service at the backend, and it talks about configuring auth for the "azure resource" instead of HTTP(S), it's an automatic NO.

upvoted 5 times

You are developing a .NET Core MVC application that allows customers to research independent holiday accommodation providers.

You want to implement Azure Search to allow the application to search the index by using various criteria to locate documents related to accommodation venues.

You want the application to list holiday accommodation venues that fall within a specific price range and are within a specified distance to an airport.

What should you do?

- A. Configure the SearchMode property of the SearchParameters class.
- B. Configure the QueryType property of the SearchParameters class.
- C. Configure the Facets property of the SearchParameters class.
- D. Configure the Filter property of the SearchParameters class.

**Suggested Answer: D**

The Filter property gets or sets the OData \$filter expression to apply to the search query.

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters> <https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.querytype>

Community vote distribution

D (100%)

🗳️ 👤 **7ack** Highly Voted 2 years, 6 months ago

Azure search was in AZ-203, but it's not in the AZ-204 list of measured skills. Is this even still relevant?

upvoted 21 times

🗳️ 👤 **Bogdan75** 1 year, 9 months ago

You're right, it's not on the list of measured skills: <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B>

upvoted 3 times

🗳️ 👤 **hstml** Highly Voted 2 years, 3 months ago

The solution D - Filter is correct.

The thing is that you filter the queried information by a price range and distance to airport. The QueryType is about the syntax of what you can pass into the search.

Read the explanations in the docs: <https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters?view=azure-dotnet>

upvoted 13 times

🗳️ 👤 **MiraA** 2 years, 2 months ago

Geo-search example:

<https://docs.microsoft.com/en-us/azure/search/search-query-simple-examples#example-6-geo-search>

upvoted 4 times

🗳️ 👤 **PhilLI** 1 year, 11 months ago

So [A] should be the answer then?

<https://docs.microsoft.com/en-us/azure/search/query-simple-syntax>

The searchMode parameter is relevant in this example. Whenever boolean operators are on the query, you should generally set searchMode=all to ensure that all of the criteria is matched. Otherwise, you can use the default searchMode=any that favors recall over precision.

upvoted 1 times

🗳️ 👤 **PhilLI** 1 year, 11 months ago

I see I'm wrong ...

upvoted 1 times

🗳️ 👤 **aruni\_mishra** 1 year, 11 months ago

correct link:

<https://docs.microsoft.com/en-us/azure/search/search-query-simple-examples#example-6-geospatial-search>

upvoted 1 times

🗄️ 👤 **vizay** Most Recent 2 months, 1 week ago

**Selected Answer: D**

- ☐ SearchMode: Chooses between matching any or all keywords — not for regex.
- ☐ QueryType: Enables advanced search features like regex when set to "full". ✓
- ☐ Facets: Groups and counts results — used for filters like "by city."
- ☐ Filter: Applies exact conditions (e.g., price < 100) — not for pattern matching.
- ✓ Use QueryType to allow regex in Azure Search.

upvoted 1 times

🗄️ 👤 **forgetfulalligator** 10 months, 1 week ago

**Selected Answer: D**

The Filter property is used to narrow down search results by specifying one or more filter expressions to apply to the search query. A filter expression is a Boolean expression that is used to match specific criteria, such as a specific price range or proximity to an airport.

upvoted 2 times

🗄️ 👤 **wrak** 10 months, 2 weeks ago

out-of-the-scope for the AZ-204 Exam; please remove from the Question Bank

upvoted 5 times

🗄️ 👤 **examinator007** 11 months ago

Azure Search is not part of the AZ-204 - see: <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B>

upvoted 2 times

🗄️ 👤 **alexein74** 11 months, 1 week ago

**Selected Answer: D**

D. Configure the Filter property of the SearchParameters class.

To allow the application to list holiday accommodation venues that fall within a specific price range and are within a specified distance to an airport, you should configure the Filter property of the SearchParameters class in your .NET Core MVC application. The filter property allows you to filter the search results by applying conditions to specific fields in the indexed documents. You can use this property to specify the price range and distance range of the accommodation venues.

upvoted 1 times

🗄️ 👤 **KingChuang** 1 year ago

**Selected Answer: D**

My Answer:D

Ref:<https://learn.microsoft.com/en-us/azure/search/search-query-simple-examples#example-5-range-filters>

upvoted 1 times

🗄️ 👤 **Chang401** 1 year, 2 months ago

**Selected Answer: D**

similar type of question is there on examtopics az-204. but the ask is different in both the questions.

upvoted 2 times

🗄️ 👤 **PhilLI** 1 year, 11 months ago

**Selected Answer: D**

checkout MiraA's link

upvoted 3 times

🗄️ 👤 **ning** 2 years, 3 months ago

querytype cannot work for geo search, it is D, filter with odata can work with geo data

upvoted 2 times

🗄️ 👤 **Bishopel** 2 years, 4 months ago

[https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.filter?view=azure-dotnet#Microsoft\\_Azure\\_Search\\_Models\\_SearchParameters\\_Filter](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.models.searchparameters.filter?view=azure-dotnet#Microsoft_Azure_Search_Models_SearchParameters_Filter)

upvoted 1 times

🗄️ 👤 **Bishopel** 2 years, 4 months ago



Correct answer is D.

upvoted 3 times

🗄️ 👤 **hussamAlHunaiti** 2 years, 4 months ago

Wrong answer, the correct answer is B- Querytype property

upvoted 2 times

  **wolf\_lu** 2 years, 5 months ago

choose B ,querytype

upvoted 2 times

You are a developer at your company.  
You need to edit the workflows for an existing Logic App.  
What should you use?

- A. the Enterprise Integration Pack (EIP)
- B. the Logic App Code View
- C. the API Connections
- D. the Logic Apps Designer


**Suggested Answer: A**

For business-to-business (B2B) solutions and seamless communication between organizations, you can build automated scalable enterprise integration workflows by using the Enterprise Integration Pack (EIP) with Azure Logic Apps.

Reference:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-overview> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions>

Community vote distribution



**Nokaido** Highly Voted 2 years, 3 months ago

**Selected Answer: D**

For definitions use the Code View, for the Workflows use the Designer I think  
upvoted 49 times

**Ming0808** 2 years, 1 month ago

yes. Test in Azure  
upvoted 4 times

**ryvel** Highly Voted 2 years, 7 months ago

Since the question says developer, I think B(the Logic App Code View) should be the answer  
upvoted 12 times

**kishoremnk** Most Recent 10 months, 1 week ago

for code changes in json use code view, for designing workflows use app designer, to edit workflow use EIP  
upvoted 2 times

**Isoldhe** 10 months, 2 weeks ago

**Selected Answer: D**

The most appropriate option for editing workflows in an existing Logic App would be:

D. the Logic Apps Designer

The Logic Apps Designer provides a visual interface for creating and modifying Logic App workflows. It allows you to easily design workflows using pre-built connectors and triggers without needing to write code. This option is commonly used for making changes to existing Logic App workflows.  
upvoted 3 times

**manopeydakon** 11 months, 2 weeks ago

D. The Logic Apps Designer is a visual interface provided by Azure Logic Apps that allows you to design, modify, and manage your logic app workflows.  
upvoted 1 times

**[Removed]** 11 months, 2 weeks ago

**Selected Answer: D**

Gpt sad that is D  
upvoted 2 times

**manopeydakon** 12 months ago

Correct answer is: The Logic Apps Designer is the graphical interface provided by Azure Logic Apps for designing and editing workflows. It allows you to visually create, edit, and manage the workflow by connecting various actions and triggers.

Option B (Logic App Code View) is used when you want to view or edit the underlying workflow definition in JSON format. If you prefer to work with the code or need to make specific changes that are easier to express in JSON, you can use the Code View.

upvoted 2 times

🗳️ 👤 **Ciupaz** 1 year, 1 month ago

**Selected Answer: B**

The Logic App Code View is the preferred method for editing Logic App workflows. It provides a JSON-based editor that allows developers to have full control over the workflow definition. This includes the ability to add, modify, and delete workflow steps, as well as configure the properties of each step.

upvoted 1 times

🗳️ 👤 **Ciupaz** 1 year, 1 month ago

Anyway, the Logic Apps Designer could be a good choice too.

upvoted 1 times

🗳️ 👤 **dddddd111** 1 year, 2 months ago

chatGPT chose B and D.

You can indeed use both the Logic App Code View (option B) and the Logic Apps Designer (option D) to edit workflows for an existing Logic App. The choice between them depends on your preference and the specific requirements of your task.

Logic App Code View (option B): This allows you to edit the Logic App's underlying JSON definition directly. It provides more granular control and is suitable for developers who prefer working with JSON or need to make intricate changes to the Logic App's structure.

Logic Apps Designer (option D): This provides a visual interface to design Logic Apps using pre-built triggers, actions, and connectors. It's a more user-friendly option for those who prefer a graphical approach and don't need to work directly with the JSON definition.

Both options are valid tools for editing workflows, and the choice depends on your familiarity with JSON, your specific task requirements, and personal preference.

upvoted 1 times

🗳️ 👤 **garbas** 1 year, 2 months ago

**Selected Answer: D**

Both Code View and Designer can be used to edit the workflow.

But it's easier to edit the workflow in the designer, so I think that is what we *should* use.

We shouldn't use the more complicated tool just because we are developers.

upvoted 3 times

🗳️ 👤 **HimGarg** 1 year, 2 months ago

Edit Workflow, so Designer

upvoted 3 times

🗳️ 👤 **heatfan900** 1 year, 3 months ago

D. the Logic Apps Designer

upvoted 2 times

🗳️ 👤 **RikinPatel** 1 year, 3 months ago

May be "Enterprise Integration Pack (EIP)" correct ans as per this portal because we don't know context of this que. in exam may be given some case study before this que.

upvoted 1 times

🗳️ 👤 **sen\_su** 1 year, 3 months ago

How to understand the given explanation "For business-to-business (B2B) solutions and seamless communication between organizations, you can build automated scalable enterprise integration workflows by using the Enterprise Integration Pack (EIP) with Azure Logic Apps." to edit the workflow for the logic App?

upvoted 1 times

🗳️ 👤 **CellCS** 1 year, 5 months ago

Logic App Code View because you are developer. designer is for "designer" who do not need has coding skills



upvoted 1 times

🗳️ 👤 **tiagodfsouza** 1 year, 5 months ago

**Selected Answer: A**

Should be A

upvoted 1 times

  **viking2** 1 year, 6 months ago

**Selected Answer: D**

There is nothing in the question about EIP.  
I'm a developer and I still use the Designer.  
I'd answer D

upvoted 3 times

## DRAG DROP -

You are a developer for a company that provides a bookings management service in the tourism industry. You are implementing Azure Search for the tour agencies listed in your company's solution.

You create the index in Azure Search. You now need to use the Azure Search .NET SDK to import the relevant data into the Azure Search service.

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

Select and Place:

**Answer Area**

Create a DataSource instance and set its Container property to the DataContainer.

Create an IndexBatch that contains the documents which must be added.

Set the DataSources property of the SearchServiceClient.

Create a SearchIndexClient object to connect to the search index.

Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.

**Suggested Answer:****Answer Area**

Create a DataSource instance and set its Container property to the DataContainer.

Set the DataSources property of the SearchServiceClient.

Call the Documents.Suggest method of the SearchIndexClient and pass the DataSource.

Create a SearchIndexClient object to connect to the search index.

Create an IndexBatch that contains the documents which must be added.

Call the Documents.Index method of the SearchIndexClient and pass the IndexBatch.

1. The index needs to be populated. To do this, we will need a SearchIndexClient. There are two ways to obtain one: by constructing it, or by calling

Indexes.GetClient on the SearchServiceClient. Here we will use the first method.

2. Create the indexBatch with the documents



Something like:

```
var hotels = new Hotel[];
```

```
{
```

```
new Hotel()
```

```
{
```

```
HotelId = "3",
```

```
BaseRate = 129.99,
```

```
Description = "Close to town hall and the river"
```

```
}
```

```
};
```

```
};
```

```
var batch = IndexBatch.Upload(hotels);
```

3. The next step is to populate the newly-created index

Example:

```
var batch = IndexBatch.Upload(hotels);
```

```
try
```

```
{
```

```
indexClient.Documents.Index(batch);
```

```
}
```

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk>

  **[Removed]**  2 years, 9 months ago

Azure search is out of scope for AZ-204

upvoted 32 times

  **wrak**  2 years, 4 months ago

out-of-the-scope for the AZ-204 Exam; please remove from the Question Bank

upvoted 9 times

  **aks\_exam**  8 months, 2 weeks ago

Azure Search has been replaced by Azure Cognitive Search.

It means not that this Q is out of scope

upvoted 1 times

  **Zinat** 1 year, 9 months ago

azure search is deprecated

upvoted 4 times

  **moota** 2 years, 1 month ago

It's also possible the indexClient comes after creating the batch.

```
// Assume that you have a collection of documents called "documents"
```

```
List<MyDocumentType> documents = GetDocuments(); // Replace MyDocumentType with the actual type of your documents
```

```
// 1. Create an IndexBatch that contains the documents which must be added.
```

```
IndexBatch<MyDocumentType> batch = IndexBatch.Upload(documents);
```

```
// 2. Create a SearchIndexClient object to connect to the search index.
```

```
// Assume that searchIndexName is the name of your Azure Search index
```

```
SearchIndexClient indexClient = serviceClient.Indexes.GetClient(searchIndexName);
```

upvoted 1 times

  **tom112** 2 years, 1 month ago

it's SearchClient.IndexDocuments(batch) method, actually.

upvoted 1 times

  **NombreFalso** 2 years, 4 months ago

Bad question

upvoted 1 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

Looks correct

<https://learn.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk#create-an-index>

upvoted 2 times

🗨️ 👤 **Discovery\_** 2 years, 7 months ago

Azure Search doesn't seem to be in the latest study guide. Best to avoid this question for less cram.

upvoted 4 times

🗨️ 👤 **coffecold** 2 years, 9 months ago

Can't be true : uploading the data by putting it hardcoded in the source? Would suggest the solutions that use a datasource. After all you had to extract the data from the "company's solution"

upvoted 2 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

It's a simple prototype. TDD, baby.

upvoted 1 times

🗨️ 👤 **sghaha** 3 years, 2 months ago

<https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk>

upvoted 2 times

You are developing an application that applies a set of governance policies for internal and external services, as well as for applications. You develop a stateful ASP.NET Core 2.1 web application named PolicyApp and deploy it to an Azure App Service Web App. The PolicyApp reacts to events from

Azure Event Grid and performs policy actions based on those events.

You have the following requirements:

- ⇒ Authentication events must be used to monitor users when they sign in and sign out.
- ⇒ All authentication events must be processed by PolicyApp.
- ⇒ Sign outs must be processed as fast as possible.

What should you do?


- A. Create a new Azure Event Grid subscription for all authentication events. Use the subscription to process sign-out events.
- B. Create a separate Azure Event Grid handler for sign-in and sign-out events.
- C. Create separate Azure Event Grid topics and subscriptions for sign-in and sign-out events.
- D. Add a subject prefix to sign-out events. Create an Azure Event Grid subscription. Configure the subscription to use the subjectBeginsWith filter.

**Suggested Answer: D**

Reference:

<https://docs.microsoft.com/en-us/azure/event-grid/subscription-creation-schema>

Community vote distribution



ning Highly Voted 3 years, 10 months ago

Only C is mentioned both topic and subscription, which are two critical parts for event grid, so I will go C, anyway, this question is very poorly worded  
upvoted 38 times

Bishopel Highly Voted 3 years, 10 months ago

For your interest, please see <https://docs.microsoft.com/en-us/azure/event-grid/event-filtering>

Correct answer is D.

upvoted 17 times

ucsdmiami2020 3 years, 6 months ago

I agree with D per the following text found in Microsoft docs URL <https://docs.microsoft.com/en-us/azure/event-grid/concepts>

"A subscription tells Event Grid which events on a topic you're interested in receiving. When creating the subscription, you provide an endpoint for handling the event. You can filter the events that are sent to the endpoint. You can filter by event type or subject pattern."

upvoted 2 times

edengoforit 3 years, 5 months ago

D doesn't contain any action of Sign-in. I am more inclined to C

upvoted 9 times

prode Most Recent 4 months, 1 week ago

**Selected Answer: D**

The requirement states that sign-out events must be processed as fast as possible while ensuring that all authentication events are processed by PolicyApp.

Azure Event Grid allows filtering of events based on subject prefixes using the subjectBeginsWith filter.

By adding a subject prefix to sign-out events (e.g., /auth/signout), you can create a subscription that only processes sign-out events separately, ensuring faster processing.

Sign-in events can be handled by a separate general subscription that processes all authentication events.

upvoted 1 times

VK1989 5 months, 3 weeks ago

**Selected Answer: D**

Answer should be D, to cross verify I used Copilot and Gemini and both accept D

upvoted 1 times

🗳️ 👤 **zcmnt** 8 months, 1 week ago

**Selected Answer: D**

D

upvoted 1 times

🗳️ 👤 **LongLasting** 8 months, 3 weeks ago

**Selected Answer: C**

Since sign-out events need to be processed as quickly as possible, having a dedicated topic and subscription for them ensures they are not delayed by sign-in events or other types of events.

upvoted 2 times

🗳️ 👤 **Zidano** 9 months, 2 weeks ago

Answer is D. All authentication events are processed by PolicyApp.

By adding a subject prefix to unsubscribe events and using the filter, unsubscribe events can be quickly identified and prioritized.

Meets both requirements:

All events are processed by PolicyApp.

Unsubscribe events are processed as quickly as possible.

upvoted 1 times

🗳️ 👤 **8ac3742** 11 months ago

The answer is correct, Option B creates a separate Azure subscription for sign-in and sign-out events implying that two event types are required, subject prefix just require one event type for sign-in and sign-out events.

upvoted 2 times

🗳️ 👤 **jobolesonihal** 1 year, 3 months ago

**Selected Answer: C**

First, the answers are poorly written.

To process sign-outs faster, we will have to do parallel processing. (there is no priority assignment in Event Grid)

Ideal solution would be to create one topic and 2 different subscriptions -

a) without any filter

b) another with just sign-out filter

Option D has 2 issues - a) It expects change to the external app i.e. Azure App Service Web App (remember you are developing PolicyApp not the app being monitored) to log the sign-out event with a prefix.

b) It only takes care of Sign-out events. To cater to all events, you would need yet another subscription.

I don't like C either. One doesn't need different Topic as it could also mean impact to the application ingesting these events to split the events to 2 different topics. C is inefficient way to solve the problem.

upvoted 2 times

🗳️ 👤 **FeriAZ** 1 year, 4 months ago

**Selected Answer: C**

By creating dedicated topics (and corresponding subscriptions) for sign-out events, you ensure that these events can be processed immediately and without the overhead of filtering from a mixed set of authentication events. This approach provides a clear and direct path for sign-out events to be handled, potentially reducing processing time and meeting the prioritization requirement.

upvoted 3 times

🗳️ 👤 **CarlosTheBoldest** 1 year, 7 months ago

**Selected Answer: D**

The "All authentication events must be processed by PolicyApp clause" should avoid C, so is D

upvoted 1 times

🗳️ 👤 **Mick1z8** 1 year, 8 months ago

**Selected Answer: C**

Filtering is something that you can use regardless (which adds a small overhead though). But the question says "sign-outs should be as fast as possible" which means creating a dedicated topic only for "sign-outs" is the best possible solution.

upvoted 2 times

🗳️ 👤 **INGAKULA** 1 year, 9 months ago

The correct answer is C

upvoted 1 times

🗨️ 👤 **MVCChicago** 1 year, 10 months ago

Answer is D

upvoted 2 times

🗨️ 👤 **LxstlnNxght** 1 year, 10 months ago

**Selected Answer: C**

It's C

upvoted 1 times

🗨️ 👤 **jakobste** 1 year, 10 months ago

I think i have seen this before and the answer was C on that practice test.

upvoted 1 times

🗨️ 👤 **Vaibhab** 1 year, 11 months ago

Answer is D:

By adding a subject prefix to sign-out events, you can differentiate them from other authentication events. Then, you can create an Azure Event Grid subscription and configure it to use the subjectBeginsWith filter. This filter allows you to specify that you only want to process events with a subject that begins with the specified prefix, in this case, the sign-out events.

upvoted 1 times

## HOTSPOT -

You are developing a C++ application that compiles to a native application named process.exe. The application accepts images as input and returns images in one of the following image formats: GIF, PNG, or JPEG.

You must deploy the application as an Azure Function.

You need to configure the function and host json files.

How should you complete the json files? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

function.json

{

	▼
"type": "http"	
"platform": "gcm"	
"datatype": "stream"	
"path": "process.exe"	

"direction": "out",  
"name" : "result"

}

host.json

	▼
"customHandler": { "description": {	
"languageWorker": { "path": {	
"extensions": { "worker": {	
"extensionBundle": {	

"defaultExecutablePath": "process.exe"

},

	▼
"enableForwardingHttpRequest": true	
"enableForwardingHttpRequest": false	

}

```
function.json
```

```
{  
  {  
    "type": "http"  
    "platform": "gcm"  
    "datatype": "stream"  
    "path": "process.exe"  
  }  
  "direction": "out",  
  "name" : "result"  
}
```

```
host.json
```

Suggested Answer:

```
{  
  {  
    "customHandler": { "description": {  
      "languageWorker": { "path": {  
        "extensions": { "worker": {  
          "extensionBundle": {  
            "defaultExecutablePath": "process.exe"  
          },  
          {  
            "enableForwardingHttpRequest": true  
            "enableForwardingHttpRequest": false  
          }  
        }  
      }  
    }  
  }  
}
```

Box 1: "type": "http"

Box 2: "customHandler": { "description": {

A custom handler is defined by configuring the host.json file with details on how to run the web server via the customHandler section.

The customHandler section points to a target as defined by the defaultExecutablePath.

Example:

```
"customHandler": {  
  "description": {  
    "defaultExecutablePath": "handler.exe"
```

Box 3: "enableForwardingHttpRequest": false

Incorrect:

For HTTP-triggered functions with no additional bindings or outputs, you may want your handler to work directly with the HTTP request and response instead of the custom handler request and response payloads. This behavior can be configured in host.json using the enableForwardingHttpRequest setting.

At the root of the app, the host.json file is configured to run handler.exe and enableForwardingHttpRequest is set to true.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers>

 **[Removed]**  2 years, 9 months ago

I just hope the real questions are not this deep to know if enableForwardingHttpRequest is true or false. I use Azure functions every day but I never used a custom handler and I hope most people do not have the burning need to use this when there are so many other solutions to handle this.  
upvoted 35 times

 **Esward**  2 years, 5 months ago

HTTP, Custom handler, true are the correct answers  
upvoted 11 times

 **Ody** 6 months, 1 week ago

The `enableForwardingHttpRequest`` setting in Azure Functions custom handlers determines how the HTTP request is forwarded to your custom handler.

### Example Use Cases

`true` - If you have a custom handler written in a language like Go or C++ that expects to handle HTTP requests directly, you would set `enableForwardingHttpRequest`` to `true``.

`false` - If your custom handler is designed to process structured data and you want the Azure Functions runtime to handle the HTTP specifics, you would set it to `false``.

<https://github.com/Azure/azure-functions-host/issues/6462>: [Azure Functions custom handlers](https://learn.microsoft.com/en-us/azure/azure-functions/functions/custom-handlers)

upvoted 3 times

  **prode**  4 months, 1 week ago

"enableForwardingHttpRequest": true

Why?

Since the C++ application (process.exe) is a native executable, it cannot directly handle HTTP requests.

Azure Functions must forward HTTP requests to the custom handler (C++ app).

Setting `enableForwardingHttpRequest: true` ensures that Azure Functions acts as a proxy, forwarding requests to process.exe for processing.

When would `enableForwardingHttpRequest: false` be used?

If the function does not depend on HTTP request details (e.g., input is processed from a queue or blob storage).

If the function retrieves input from other bindings instead of HTTP.


Since the question clearly states that the function receives images and returns processed images, HTTP request forwarding is required.

upvoted 1 times

  **8ac3742** 11 months ago

"enableForwardingHttpRequest" should be set to true to forward the HTTP request to your custom handler.

upvoted 2 times

  **oskx2** 1 year, 2 months ago

`enableForwardingHttpRequest` must be true: Ref: <https://learn.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other?tabs=go%2Cwindows>

In the customHandler section, add a property named `enableForwardingHttpRequest` and set its value to true. For functions consisting of only an HTTP trigger, this setting simplifies programming by allow you to work with a typical HTTP request instead of the custom handler request payload.

upvoted 2 times

  **manopeydakon** 1 year, 5 months ago

If your C++ application Azure Function does not depend on information from the original HTTP request (like headers, query parameters, etc.), you might set `enableForwardingHttpRequest` to false to optimize performance and reduce unnecessary data transfer.

If your C++ application Azure Function requires information from the original HTTP request, such as headers or query parameters, you might set `enableForwardingHttpRequest` to true so that the original request is forwarded to your function.

So I would say True for this case!

upvoted 3 times

  **Stel0Papad4** 1 year, 1 month ago

I dont think its `{enableForwardingHttpRequest:true}` cause it says that it receives images and returns images in a specified format. Input image could differ from the output one. The question is very tricky and bad written but it doesnt say that the the given image is returned back. So my understanding is that the C++ App Azure Function doesnt depend on information from the original HTTP Request.

Correct me if im wrong.

upvoted 2 times

  **130nk3r5** 1 year, 6 months ago

Got this today.

Went with answer here.

Score 927

upvoted 8 times

  **dddddd111** 1 year, 7 months ago



It should be "enableForwardingHttpRequest": true  
upvoted 3 times

🗨️ 👤 **ENGs** 1 year, 8 months ago

On my exam 2023-10 before the Update of the Exam  
upvoted 2 times

🗨️ 👤 **jakobste** 1 year, 10 months ago

This was on my exam in august 2023.  
upvoted 8 times

🗨️ 👤 **AbidooKing** 1 year, 11 months ago

On my exam 2023-07-13- went with "enableForwardingHttpRequest": true. Score 917  
upvoted 5 times

🗨️ 👤 **Gednehm** 1 year, 12 months ago

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers#>, the last one should be true.  
upvoted 3 times

🗨️ 👤 **moota** 2 years, 1 month ago

- Your Azure Function will be a custom handler since Azure Functions doesn't directly support C++ as a runtime.  
- enableForwardingHttpRequest is set to true to allow the custom handler to receive the HTTP request data  
upvoted 6 times

🗨️ 👤 **surprise0011** 2 years, 2 months ago

I will go with 'yes'

In the function.json file, enableForwardingHttpRequest should be set to true for the Azure Function to accept HTTP requests and forward them to the C++ application for processing.

This is because the C++ application is compiled as a native application and cannot directly receive HTTP requests. By setting enableForwardingHttpRequest to true, the Azure Function acts as a proxy and forwards incoming HTTP requests to the C++ application.

So the correct option is:

"enableForwardingHttpRequest": true  
upvoted 5 times

🗨️ 👤 **Videira** 2 years, 4 months ago

On my exam 2023-02-25  
upvoted 7 times

🗨️ 👤 **forgetfulalligator** 2 years, 4 months ago

Everything looks gucci except enableforwardhttprequest got to be true. If it was false, our customer handler can receiving any incoming HTTP requests  
upvoted 2 times

🗨️ 👤 **AAcademy** 2 years, 4 months ago

enableForwardingHttpRequest should be true

When enableForwardingHttpRequest is true, the behavior of HTTP-only functions differs from the default custom handlers behavior in these ways:

The HTTP request does not contain the custom handlers request payload. Instead, the Functions host invokes the handler with a copy of the original HTTP request.

The Functions host invokes the handler with the same path as the original request including any query string parameters.

The Functions host returns a copy of the handler's HTTP response as the response to the original request.

Source: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers>

upvoted 6 times

## HOTSPOT

-

You are developing an Azure Static Web app that contains training materials for a tool company. Each tool's training material is contained in a static web page that is linked from the tool's publicly available description page.

A user must be authenticated using Azure AD prior to viewing training.

You need to ensure that the user can view training material pages after authentication.

How should you complete the configuration file? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```
{
  "routes": {
    "routes":
    headers
    responseOverrides
    navigationFallback
  },
  "statusCode": 302
  "redirect": "/auth/login/ ?post_id"
  "status": {
    400
    401
    403
    404
  },
  "status": {
    add
    azure
    graph
    microsoftonline
  },
  "status": {
    orig
    route
    return
    .referrer
  }
}
```

## Answer Area

```
{
  "routes": {
    "headers": {
      "responseOverrides": {
        "navigationFallback": {
          "statusCode": 302,
          "redirect": "/auth/login/?post_login_redirect_uri=.referrer"
        }
      }
    }
  }
}
```

Suggested Answer:

 **[Removed]** Highly Voted 1 year, 9 months ago

Got it in the exam 7/03/23. Went with the given answer  
upvoted 9 times

 **chettir01** Highly Voted 1 year, 11 months ago


That's correct :  
<https://learn.microsoft.com/en-us/azure/static-web-apps/configuration>  
upvoted 9 times

 **JustHereToLearn** 1 year, 11 months ago

How that answer is correct? The first two are okay but add does not make sense it should be "aad" see reference <https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#restrict-access-to-entire-application>  
Also not sure what the last answer is but it is surely not referrer as there no such thing I could find in the docs.  
upvoted 5 times

 **warchoon** 1 year, 10 months ago

looks like after scan typo  
upvoted 5 times

 **alejary** 1 year, 11 months ago

I think this should look like this (.referrer will be replaced with the URL of original request):

```
"responseOverrides": {
  "401": {
    "statusCode": "302",
    "redirect": "/.auth/login/aad?post_login_redirect_uri=.referrer"
  }
}
```

upvoted 38 times

 **Mirdwal** Most Recent 1 month, 3 weeks ago

Top Dropdown – responseOverrides  
This section is used to override responses for certain HTTP status codes (like 401 Unauthorized).

Next Dropdown – 401

You want to catch 401 Unauthorized errors, which occur when a user isn't authenticated, and redirect them to login.



Redirect Path – /.auth/login/aad

This is the correct path to redirect users to Azure Active Directory login for Static Web Apps.

Status Code – 302



This means "Found", and it tells the browser to redirect the user to the login page.

upvoted 2 times

  **onlyforheros** 9 months, 2 weeks ago

Got it in the exam 13.03.2024. Score: 910. Went with the give answer

upvoted 6 times

  **egaws** 10 months, 3 weeks ago

Overrides

Aad

redirect

<https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#restrict-access-to-entire-application>

upvoted 1 times

  **Ciupaz** 12 months ago

HTTP 401 -> Unauthorized

upvoted 4 times

  **bgbgvfvf** 1 year ago

answer is correct

<https://learn.microsoft.com/en-us/azure/static-web-apps/configuration>

upvoted 1 times

  **trueinam** 1 year ago

This question was on exam 2023-12-21.

upvoted 3 times

  **arunkuml** 1 year ago

Got it in the exam 14/12/23. Went with given answer. All questions are from ExamTopics. Case study - VanArsdel, Ltd (11 questions)

upvoted 2 times

  **dddddd111** 1 year, 1 month ago

FYI Azure Active Directory (aad) is replaced with MS Entra ID as of Nov14,2023.



upvoted 2 times

  **dddddd111** 1 year, 1 month ago

Acronym usage is not encouraged, but if you must replace AAD with an acronym due to space limitations, use ME-ID.

<https://learn.microsoft.com/en-us/entra/fundamentals/new-name>

upvoted 1 times

  **p2006** 1 year, 3 months ago

Got on 9/25/2023



responseoverrides

401

aad

.referrer

upvoted 5 times

  **RuffBoii** 1 year, 3 months ago

Had this on my exam today. Went with the given answer.

upvoted 1 times

  **mihailos** 1 year, 4 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 2 times

  **Masood07** 1 year, 1 month ago

How do you all copy during the exam. please mention.

upvoted 2 times

  **ReyPirata** 1 year, 4 months ago

This was on the exam (08/20/2023). Went with proposed. Scored 925  
upvoted 2 times

🗲️ 👤 **juanckar** 1 year, 5 months ago

This was on the exam (July 2023). Went with proposed. Scored 917  
upvoted 4 times

🗲️ 👤 **deathRac3** 1 year, 7 months ago

Question was there for me on 29th May 2023  
upvoted 4 times

🗲️ 👤 **narenazure** 1 year, 8 months ago

<https://learn.microsoft.com/en-us/azure/static-web-apps/configuration#restrict-access-to-entire-application>

```
"responseOverrides": {  
  "401": {  
    "statusCode": 302,  
    "redirect": "/.auth/login/aad"  
  }  
}
```

upvoted 4 times

## HOTSPOT

-

You are authoring a set of nested Azure Resource Manager templates to deploy Azure resources. You author an Azure Resource Manager template named `mainTemplate.json` that contains the following linked templates: `linkedTemplate1.json`, `linkedTemplate2.json`.

You add parameters to a parameters template file named `mainTemplate.parameters.json`. You save all templates on a local device in the `C:\templates\` folder.

You have the following requirements:



- Store the templates in Azure for later deployment.
- Enable versioning of the templates.
- Manage access to the templates by using Azure RBAC.
- Ensure that users have read-only access to the templates.
- Allow users to deploy the templates.

You need to store the templates in Azure.

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

NOTE: Each correct selection is worth one point.


**Answer Area**

```
az  \  
ts create  
storage account create  
storage account update  
blueprint artifact template create  
  
--name templateStore \  
--version "1.0" \  
--resource-group templatesRG \  
--location "eastus" \  
  
--template-file  "\  
mainTemplate.json  
linkedTemplate1.json  
linkedTemplate2.json  
mainTemplate.parameters.json  
  
"C:\templates\  
--tags Dept=HumanResources Environment=Production
```

## Answer Area

Suggested Answer:


```
az  \
--name templateStore \
--version "1.0" \
--resource-group templatesRG \
--location "eastus" \
--template-file  "\
"C:\templates\
--tags Dept=HumanResources Environment=Production
```

 **narenazure** Highly Voted 2 years, 5 months ago

correct

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/template-specs-create-linked?tabs=azure-cli>

upvoted 16 times

 **Esward** 2 years, 5 months ago

I agree

upvoted 1 times

 **Dotar** 2 years, 5 months ago


I agree the second option is because it needs the path to main template no to the parameters.

upvoted 2 times

 **ReyPirata** Highly Voted 1 year, 10 months ago

This was on the exam (08/20/2023). Went with proposed. Scored 925

upvoted 11 times

 **TheFivePips** Most Recent 10 months, 2 weeks ago

The `ts create` command creates a Template Spec. Template Specs in Azure allow you to create, manage, and reuse Azure Resource Manager (ARM) templates as reusable artifacts. They provide a way to standardize and streamline the deployment of Azure resources across your organization by encapsulating ARM templates in a versioned, reusable format.

upvoted 3 times

 **1CY1** 1 year ago

Deploy template spec

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-cli#deploy-template-spec>

Instead of deploying a local or remote template, you can create a template spec.

The template spec is a resource in your Azure subscription that contains an ARM template.

```
az ts create --name storageSpec --version "1.0"
```

```
--resource-group templateSpecRG --location "westus2"
```

```
--template-file "./mainTemplate.json"
```

upvoted 1 times

 **[Removed]** 1 year, 11 months ago

Why not use `az storage account create`?


We need to create a storage account first before mentioning the name of the template file..

upvoted 1 times

 **macobuzi** 1 year, 10 months ago

Because the code is not about creating a storage account. Pay attention to the part: `--name TemplateStore`

upvoted 6 times

 **viking2** 1 year, 12 months ago

Correct.

The parameters given to the command match those of the az ts command, but not those of az blueprint, so the "ts create" and the main template file are correct answers.

upvoted 2 times

  **nedlo** 2 years, 3 months ago



No, the az ts command is not for managing Azure Blueprints.

The az ts command is used for managing Azure Time Series Insights, which is a service that allows you to store, visualize, and query large amounts of time-series data in the cloud. With az ts command, you can create and manage Azure Time Series Insights environments, as well as interact with the data stored in those environments.

Azure Blueprints, on the other hand, is a service that allows you to define a repeatable set of Azure resources that adhere to organizational standards, patterns, and requirements. With Azure Blueprints, you can create a blueprint that specifies the resources and configurations needed to meet your organization's standards and deploy it to multiple subscriptions and environments.

To manage Azure Blueprints, you can use the Azure portal, Azure PowerShell, Azure CLI, or Azure REST API. The Azure CLI command for managing Azure Blueprints is az blueprint, which allows you to create, manage, and assign blueprints to Azure environments.

upvoted 1 times

  **rgullini** 2 years, 2 months ago

This whole text is wrong. Please, ignore this answer

upvoted 17 times

  **Mattt** 2 years, 3 months ago

It looks like the wrong generated answer by ChatGPT

upvoted 3 times



  **SprintingLlama** 2 years, 3 months ago

This Microsoft article contradicts your point

<https://learn.microsoft.com/en-us/cli/azure/ts?view=azure-cli-latest>

'az ts' allows you to manage template specs at subscription or resource group scope.

upvoted 4 times

  **ucskips** 2 years, 3 months ago

Just a question. Could account storage have been used instead?

upvoted 1 times

  **chettir01** 2 years, 5 months ago

correct

<https://learn.microsoft.com/en-us/cli/azure/ts?view=azure-cli-latest>

upvoted 1 times



## HOTSPOT

-

You are developing a service where customers can report news events from a browser using Azure Web PubSub. The service is implemented as an Azure Function App that uses the JSON WebSocket subprotocol to receive news events.

You need to implement the bindings for the Azure Function App.

How should you configure the binding? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```
{
  "bindings": [
    {
      "type": "
      user
      system
      message
      connected
      webPubSubTrigger
      webPubSubConnection
      ",
      "direction": "in",
      "name": "data",
      "eventName": "message",
      "eventType": "
      user
      system
      message
      connected
      webPubSubTrigger
      webPubSubConnection
      "
    }
  ]
}
```

**Answer Area**

```
{
  "bindings": [
    {
      "type": "
      user
      system
      message
      connected
      webPubSubTrigger
      webPubSubConnection
      ",
      "direction": "in",
      "name": "data",
      "eventName": "message",
      "eventType": "
      user
      system
      message
      connected
      webPubSubTrigger
      webPubSubConnection
      "
    }
  ]
}
```

Suggested Answer:

- 🗨️ **surprise0011** Highly Voted 2 years, 2 months ago  
received 2023-17-04 went with given answer, score 926  
upvoted 24 times
- 🗨️ **adilkhan** 2 years, 2 months ago  
were all Qs from exam topics?  
upvoted 10 times
- 🗨️ **HafizSalmanMalik** Highly Voted 2 years, 5 months ago  
Given answer is correct.  
<https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript>  
upvoted 16 times
- 🗨️ **Esward** 2 years, 5 months ago  
Agreed  
upvoted 1 times
- 🗨️ **senek2137** Most Recent 10 months, 1 week ago  
you can be sure that type is user. eventName - message can only be for eventType user.  
<https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript-v3>  
upvoted 2 times
- 🗨️ **4bd3116** 10 months, 1 week ago  
Not sure, but I think the event type should be message instead of user !!  

```
{  
  "bindings": [  
    {  
      "type": "webPubSubTrigger",  
      "direction": "in",  
      "name": "data",  
      "hub": "<HUB_NAME>",  
      "eventName": "<EVENT_NAME>",  
      "eventType": "message",  
      "connectionStringSetting": "<CONNECTION_STRING_SETTING>"  
    }  
  ]  
}
```

  
upvoted 1 times
- 🗨️ **DamianSawula** 1 year, 1 month ago  
On my exam 2024-05-16, went with given answer, score 886. 44 out of 45 exam questions were exactly from here, Case Study: Contoso. I recommend learning with this site.  
upvoted 12 times
- 🗨️ **senek2137** 10 months, 1 week ago  
thanks and congrats  
upvoted 1 times
- 🗨️ **Vukile** 1 year, 7 months ago  
On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors  
upvoted 1 times
- 🗨️ **senek2137** 10 months, 1 week ago  
farmers and distributors = munson pickles. its in here  
upvoted 1 times
- 🗨️ **dddddd111** 1 year, 7 months ago  
Never seen Farmers and Distributors Case Study.  
upvoted 1 times
- 🗨️ **JojoD5** 1 year, 7 months ago  
is it a new case study?  
upvoted 1 times
- 🗨️ **AndySmith** 1 year, 7 months ago

On exam 3-Nov-2023. Went with proposed answer - 932/1000.

1) webPubSubTrigger

2) user

upvoted 2 times

🗨️ 👤 **mihailos** 1 year, 10 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 3 times

🗨️ 👤 **deathRac3** 2 years ago

Question was there for me on 29th May 2023

upvoted 4 times

🗨️ 👤 **szumik** 2 years, 4 months ago

had this question on exam today 22/02/23, answered wrongly as I saw it for the first time, but seems it's still on the syllabus

upvoted 3 times

🗨️ 👤 **red0400** 2 years, 4 months ago

pls tell me, is it right answer or not as given?

upvoted 3 times

🗨️ 👤 **adilkhan** 2 years, 2 months ago

were all Qs from exam topics?

upvoted 3 times

🗨️ 👤 **narenazure** 2 years, 5 months ago

the answer is correct but is Azure Web PubSub on the syllabus?

upvoted 2 times

🗨️ 👤 **AzSimuz** 2 years, 4 months ago

I wondered the same, I cannot find in the syllabus as of Feb 2023.

upvoted 1 times

🗨️ 👤 **pseudocode2023** 2 years, 5 months ago

Correct

<https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript#example>

upvoted 3 times

🗨️ 👤 **chettir01** 2 years, 5 months ago

Crrect

<https://learn.microsoft.com/en-us/azure/azure-web-pubsub/reference-functions-bindings?tabs=javascript>

upvoted 3 times

HOTSPOT

-

You are building a software-as-a-service (SaaS) application that analyzes DNA data that will run on Azure virtual machines (VMs) in an availability zone. The data is stored on managed disks attached to the VM. The performance of the analysis is determined by the speed of the disk attached to the VM.

You have the following requirements:

- The application must be able to quickly revert to the previous day's data if a systemic error is detected.
- The application must minimize downtime in the case of an Azure datacenter outage.

You need to provision the managed disk for the VM to maximize performance while meeting the requirements.

Which type of Azure Managed Disk should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

#### Requirement

#### Solution

Disk type

▼

Premium SSD  
Standard SSD  
Standard HDD

Redundancy

▼

Geo-redundant storage (GRS)  
Zone-redundant storage (ZRS)  
Locally-redundant storage (LRS)

### Answer Area

#### Requirement

#### Solution

Suggested Answer:

Disk type

▼

Premium SSD  
Standard SSD  
Standard HDD

Redundancy

▼

Geo-redundant storage (GRS)  
Zone-redundant storage (ZRS)  
Locally-redundant storage (LRS)

 **armpro**  2 years, 5 months ago

Wrong answer: Managed disk doesn't support GRS

Ans: Premium SSD and ZRS

They are asking for high performance workloads which is supported by Premium tier <https://learn.microsoft.com/en-us/azure/virtual-machines/disks-types>

Also they are asking for zone redundancy (if datacenter goes down, NOT region outage). Also managed disk doesn't support GRS <https://learn.microsoft.com/en-us/azure/virtual-machines/disks-redundancy>

upvoted 77 times

🗨️ 👤 **armpro** 2 years, 5 months ago

Premium SSD support backup. Premium SSD v2 doesn't support backup  
upvoted 2 times

🗨️ 👤 **8ac3742** 11 months ago

Azure Disk Backup supports multiple types of Managed Disks, including Standard HDD, Standard SSD, and Premium SSD1. However, Premium SSDs are often preferred for their superior performance and reliability, which is crucial for your DNA analysis application.  
upvoted 1 times

🗨️ 👤 **Esward** 2 years, 5 months ago

I greed. Premium SSD, ZRS are correct answers  
upvoted 6 times

🗨️ 👤 **zreaf29** 1 year, 9 months ago

I agree, too. Azure managed disk doesn't have GRS redundancy option. GRS is a option of azure storage. But this question request use azure managed disk.  
upvoted 1 times

🗨️ 👤 **narenazure** Highly Voted 🏆 2 years, 5 months ago

Azure managed disks offer two storage redundancy options, zone-redundant storage (ZRS), and locally-redundant storage.  
Ans: Premium SSD and ZRS  
<https://learn.microsoft.com/en-us/azure/virtual-machines/disks-redundancy>  
upvoted 10 times

🗨️ 👤 **jobolesonihal** Most Recent 🔔 1 year, 3 months ago

Ans: Premium SSD and ZRS  
Ref: <https://learn.microsoft.com/en-us/azure/virtual-machines/disks-redundancy>  
"Azure managed disks offer two storage redundancy options, zone-redundant storage (ZRS), and locally redundant storage. ZRS provides higher availability for managed disks than locally redundant storage (LRS) does. However, the write latency for LRS disks is better than ZRS disks because LRS disks synchronously write data to three copies in a single data center."  
upvoted 2 times

🗨️ 👤 **DaveOla** 1 year, 3 months ago

VMs are no longer covered in the exam  
upvoted 4 times

🗨️ 👤 **Razvan\_C** 1 year, 7 months ago

I believe the VMs are not in the AZ-204 anymore  
upvoted 8 times

🗨️ 👤 **heatfan900** 1 year, 9 months ago

WHEN THEY SAY 'Azure datacenter outage' they are implying a data center within a region which is hosted in one of the region's three zones. That alone gives it away that its ZRS with P-SSD.  
upvoted 2 times

🗨️ 👤 **RaoMOWais** 2 years, 3 months ago

Answer by ChatGpt :  
Premium Managed Disk  
Zone Redundant Storage (ZRS) replication.  
upvoted 1 times

🗨️ 👤 **kay000001** 1 year, 11 months ago

Or you could refer to MS docs and learn something?  
upvoted 6 times

🗨️ 👤 **Kluk\_Kluk** 2 years, 3 months ago

Premium SSD and Zone-Redundant storage ZRS  
Because Azure managed disks offer two storage redundancy options, zone-redundant storage and locally-redundant storage. ZRS provides higher availability for managed disks than LRS.  
upvoted 1 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

Maybe I am getting confused by the wording here, but doesn't 'Azure datacenter outage' mean the region is out, and thus you need GRS?  
upvoted 1 times

🗨️ 👤 **Defa** 2 years, 3 months ago

From: <https://learn.microsoft.com/en-us/azure/architecture/resiliency/recovery-loss-azure-region>

Azure is divided physically and logically into units called regions. A region consists of one or more data centers in close proximity. Many regions and services also support availability zones, which can be used to provide more resiliency against outages in a single data center. Consider using regions with availability zones to improve the availability of your solution.

\*\*\*


So, zones for outage in a datacenter.

upvoted 2 times

  **guschess** 2 years, 4 months ago

Premium and zone. Azure managed disks offer two storage redundancy options, zone-redundant storage (ZRS), and locally-redundant storage. No geo

upvoted 1 times

  **adilkhan** 2 years, 5 months ago

must be ZRS

upvoted 2 times

## HOTSPOT

-

You are developing an application that includes two Docker containers.

The application must meet the following requirements:

- The containers must not run as root.
- The containers must be deployed to Azure Container Instances by using a YAML file.
- The containers must share a lifecycle, resources, local network, and storage volume.
- The storage volume must persist through container crashes.
- The storage volume must be deployed on stop or restart of the containers.

You need to configure Azure Container Instances for the application.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area****Configuration setting    Configuration value**

Shared lifecycle

▼

Container group  
Container image  
Service endpoint  
Resource group

Storage volume

▼

Azure file share  
Secret  
Empty directory  
Cloned Git repo

**Answer Area****Configuration setting    Configuration value****Suggested Answer:**

Shared lifecycle

▼

Container group  
Container image  
Service endpoint  
Resource group

Storage volume

▼

Azure file share  
Secret  
Empty directory  
Cloned Git repo

 **armpro** Highly Voted 2 years, 5 months ago

Ans: Container group, EmptyDir

Container group is the only logical answer that can have shared lifecycle <https://learn.microsoft.com/en-us/azure/container-instances/container-instances-container-groups?source=recommendations#what-is-a-container-group>

Azure files need root permission



Secret is for secrets and read-only

EmptyDir can persist through crash and redeployed on stop and restart

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-emptydir#emptydir-volume>

Cloned Git Repo also does the job but it needs more details like Git URL and stuff which are not mentioned to be available in the question

upvoted 54 times



  **BrownieP** 5 months, 3 weeks ago

Ans: Container group, EmptyDir

Shared lifecycle: Container group. This ensures that the containers share a lifecycle, resources, local network, and storage volume.

Storage volume: Empty directory; the appropriate storage volume option would be empty directory. This option allows the containers to run without root privileges while still providing a shared storage volume that persists through container crashes

upvoted 1 times

  **TonyMel** 2 years, 3 months ago



correct, in 2023Mar24, score:904/1000

upvoted 11 times

  **narenazure** 2 years, 5 months ago

Data in an emptyDir volume is persisted through container crashes. Containers that are restarted, however, are not guaranteed to persist the data in an emptyDir volume. If you stop a container group, the emptyDir volume is not persisted.

upvoted 4 times

  **Esward** 2 years, 5 months ago


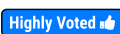
Gievn answers are correct!!

upvoted 1 times

  **alekomar** 2 years, 4 months ago



There is no requirement to persist the data in a volume between restarts

upvoted 4 times

  **GuruPratap**  2 years, 3 months ago

Chat GPT : Container Group, Azure fileshare

upvoted 10 times

  **rgullini** 2 years, 2 months ago

Cannot be Azure File Share as requirement is not to run as root.

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-azure-files>


upvoted 10 times

  **MortisG**  1 month ago

Correct answer is Container Group & Azure File Share.

Azure File Share is the only solution that provides persistent shared storage through crashes and restarts. EmptyDir does not.

upvoted 1 times

  **Thameur01** 2 months, 1 week ago

Ans: Container group, EmptyDir, please update the answer

upvoted 1 times

  **hasala** 2 months, 3 weeks ago

GPT:

Storage volume:

✓ Answer: Azure file share

□ Explanation: An Azure File Share provides persistent, durable storage that survives container crashes and restarts. This is the only option listed that fully meets the persistence requirement. Options like Empty directory are temporary and reset on restart.

upvoted 1 times

  **Kensington** 10 months, 2 weeks ago

Container group is the only logical answer, as mentioned by others.

Emptydir lifecycle is tied to the container group and therefore persists for individual container crashes.

GitRepo makes less sense as it's read-only and is mostly used for scenarios where you want to load code into your app.

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-volume-emptydir>



upvoted 1 times

🗨️ 👤 **8ac3742** 11 months ago

The Storage Volume should Azure File Share, An Empty directory (or emptyDir) in Azure Container Instances is a type of ephemeral storage volume. An Empty directory provides a writable directory that is accessible to each container in a container group. When the container group is created, the directory is initially empty, hence the name "emptyDir"1.

Empty directories are commonly used for:

scratch space: Temporary storage for data that doesn't need to persist beyond the container's lifecycle.

Checkpointing: Storing intermediate results during long-running tasks.

Sidecar containers: Sharing data between a main application container and a sidecar container, such as logs or temporary files1.

upvoted 2 times

🗨️ 👤 **Ody** 6 months, 1 week ago

I think this is the correct answer... now. Maybe not in the past.

Now, we can use Managed Identities.

```
az container create \
--resource-group <your-resource-group> \
--name <your-container-instance-name> \
--image <your-container-image> \
--assign-identity <your-managed-identity-id> \
--azure-file-volume-account-name <your-storage-account> \
--azure-file-volume-share-name <your-file-share> \
--azure-file-volume-mount-path <mount-path> \
--role StorageBlobDataContributor
```

<https://learn.microsoft.com/en-us/answers/questions/2075070/mount-volumes-in-aci-but-without-using-storage-acc>

upvoted 1 times

🗨️ 👤 **Isoldhe** 1 year, 4 months ago

ChatGPT: Azure Container Group, Azure File Share

upvoted 2 times

🗨️ 👤 **ThierryAhn** 1 year, 3 months ago

Azure file share volume mount requires the container run as root, so can not be.

upvoted 2 times

🗨️ 👤 **Christian\_garcia\_martin** 10 months, 3 weeks ago

you dont be root if you modify the ownership and permissions using the chown and chmod commands after mounting the share in your startup script or Dockerfile .So Azure File share is ok

upvoted 1 times

🗨️ 👤 **AhmedAbdelAziz** 1 year, 5 months ago

Ans: Container group, EmptyDir

EmptyDir doesn't required run as root , also support data persistence throw container crashes although it doesn't grantee persist the data in a volume between restarts as there is no requirement like that

upvoted 1 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

Container Group, Azure fileshare, You can set the user and group IDs for each container in the Docker image to ensure they do not run as the root user. This is typically done using the USER instruction in the Dockerfile. Azure File Share can be mounted as a volume in each container to provide shared storage. With these modifications, your containers will run with non-root users, and they will share storage using an Azure File Share.

upvoted 2 times

🗨️ 👤 **dddddd111** 1 year, 8 months ago

Ans: Container group and Azure File share.

upvoted 4 times

🗨️ 👤 **nayru** 1 year, 9 months ago

Got on 26/09/23 and chose container group and empty dir (score 910)


upvoted 5 times

🗨️ 👤 **katrang** 1 year, 9 months ago

Shared Lifecycle - Container Group

Storage volume - The storage volume must persist through container crashes, so emptyDir meets the requirements

upvoted 2 times

  **JH81** 2 years ago

Got this on 6/28/2023 and passed with 850. Went with Container group, EmptyDir.

upvoted 3 times

  **aragones** 2 years, 1 month ago

Got this 2023-05-12.

my cases also:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 4 times



  **GuruPratap** 2 years, 3 months ago

Container Group, Azure fileshare

For Shared lifecycle, the appropriate configuration setting would be "Container Group," which allows multiple containers to be deployed together as a group and share a local network, storage, and lifecycle.

For Storage volume, the appropriate configuration setting would be "Azure file share," which provides a durable and highly available storage solution that can be mounted as a volume in the container instance. This ensures that the storage volume persists through container crashes and can be deployed on stop or restart of the containers.

upvoted 4 times

  **rgullini** 2 years, 2 months ago

Azure File Share requires root. It is EmptyDir

upvoted 6 times

  **dddddd111** 1 year, 8 months ago

No, using Azure File Share does not require running containers as root. Azure File Share allows you to mount a shared file system to containers in Azure Container Instances without the need for elevated privileges. You can mount Azure File Share to your containers securely without running them as root, ensuring better security practices.

upvoted 3 times

  **varinder82** 2 years, 3 months ago

Could anyone give the correct answer. As it is quite confusing from the comments and multiple people having diff. options

upvoted 6 times

## HOTSPOT

-

You have a Dockerfile that builds a container image named image1. The container image and its base image are stored in separate repositories of an Azure Container registry named registry1. The codebase of image1 is stored in a GitHub repo named app1 of an account named account1.

You plan to implement automatic updates to image1 whenever its base image is updated or an update to the main branch of the GitHub repo occurs.

You need to complete the Azure CLI command that will apply the planned implementation.

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

NOTE: Each correct selection is worth one point.

## Answer Area

## Azure CLI command


az  task create --registry registry1 --name task1 --image  
  
image1:{{.Run.ID}} --arg REGISTRY\_NAME=registry1.azurecr.io  
https://github.com/account1/app1.git#main --file Dockerfile

## Suggested Answer:

Answer Area  
Azure CLI command  
az  task create --registry registry1 --name task1 --image  
  
image1:{{.Run.ID}} --arg REGISTRY\_NAME=registry1.azurecr.io  
https://github.com/account1/app1.git#main --file Dockerfile

 **Ira89** 3 months, 2 weeks ago

Agree. Correct answer  
upvoted 1 times

 **21c5438** 5 months, 1 week ago

acr  
context

<https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-base-image-update#create-a-task>

upvoted 4 times

**HOTSPOT -**

You are implementing a software as a service (SaaS) ASP.NET Core web service that will run as an Azure Web App. The web service will use an on-premises

SQL Server database for storage. The web service also includes a WebJob that processes data updates. Four customers will use the web service.

- ⇒ Each instance of the WebJob processes data for a single customer and must run as a singleton instance.
- ⇒ Each deployment must be tested by using deployment slots prior to serving production data.
- ⇒ Azure costs must be minimized.
- ⇒ Azure resources must be located in an isolated network.

You need to configure the App Service plan for the Web App.

How should you configure the App Service plan? To answer, select the appropriate settings in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area****App service plan setting****Value**

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

**Answer Area****App service plan setting****Value**

Number of VM instances

	▼
2	
4	
8	
16	

Pricing tier

	▼
Isolated	
Standard	
Premium	
Consumption	

Suggested Answer:

Number of VM instances: 4 -

You are not charged extra for deployment slots.

Pricing tier: Isolated -

The App Service Environment (ASE) is a powerful feature offering of the Azure App Service that gives network isolation and improved scale capabilities. It is essentially a deployment of the Azure App Service into a subnet of a customer's Azure Virtual Network (VNet).

Reference:

<https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/>

🗳️ 👤 **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: 4

There are four customers that use this service, and each instance of the WebJob processes data for a single customer and must run as a singleton instance. So, the number of VM should be 4. WebJobs is a feature of Azure App Service that enables you to run a program or script in the same instance as a web app. Like running background tasks.

Box 2: Isolated

Azure resources must be located in an isolated network .

In the Isolated tier, the App Service Environment defines the number of isolated workers that run your apps, and each worker is charged. In addition, there's a flat Stamp Fee for the running the App Service Environment itself. Isolated: This tier runs dedicated Azure VMs on dedicated Azure Virtual Networks. It provides network isolation on top of compute isolation to your apps. It provides the maximum scale-out capabilities.

upvoted 124 times

🗳️ 👤 **goudigubba** Highly Voted 4 years, 4 months ago

Got this in the exam yesterday.

upvoted 21 times

🗳️ 👤 **dhosth** 4 years, 3 months ago

What is the correct answer?

upvoted 2 times

🗳️ 👤 **joanbdm** 4 years, 1 month ago

4 & isolated

upvoted 8 times

🗳️ 👤 **Soumayard** 4 years ago

Heyy, did you get a lot of questions like the ones here?

upvoted 2 times

🗳️ 👤 **8ac3742** Most Recent 11 months ago

WebJob Singleton instance implies that one VM can only run one service instance at a time across all VM instances, so we require 4 VM instances.

upvoted 1 times

🗳️ 👤 **Xoxoo** 1 year, 4 months ago

Number of VM instances: Since you have four customers, each with a separate instance of the WebJob that must run as a singleton instance, you need at least four instances of the App Service plan to ensure that each customer's WebJob runs independently. Therefore, you should configure the number of VM instances to be at least four.

Pricing tier: Given the requirement for an isolated network and the need to minimize costs, the appropriate pricing tier for this scenario would be the "Isolated" tier. The Isolated tier provides dedicated infrastructure for your App Service Environment (ASE), ensuring isolation and security. While it may have higher costs compared to the Standard or Premium tiers, it offers the required level of isolation and network security for your scenario.

upvoted 1 times

🗳️ 👤 **Xoxoo** 1 year, 4 months ago

Answer: 4 and Isolated

upvoted 2 times

🗳️ 👤 **Happiman** 1 year, 4 months ago

The singleton nature of the WebJobs does not necessitate a one-to-one mapping with VM instances, as singleton execution is managed by Azure and ensures that only one instance of a job runs across the entire App Service environment. so the required VMs are two(2) not four.

upvoted 1 times

🗳️ 👤 **NPE\_** 1 year, 7 months ago

Got this on the exam Nov. 07 2023

upvoted 6 times

🗳️ 👤 **juanckar** 1 year, 11 months ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 6 times

🗳️ 👤 **Devendra\_Shukla** 1 year, 11 months ago

How many questions are from Exam topics?

upvoted 1 times

🗳️ 👤 **carlosghosn** 2 years, 7 months ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

🗨️ 👤 **drgn7676** 2 years, 11 months ago

Lorem Ipsum is simply dummy text of the printing and typesetting industry.

upvoted 3 times

🗨️ 👤 **BharatVasi** 3 years ago

Wrote exam yesterday, not even 60% question came from here :-)

upvoted 3 times

🗨️ 👤 **adilkhan** 2 years, 5 months ago

did you purchase the examtopics account and gone through all 33 pages? please reply

upvoted 3 times

🗨️ 👤 **Devendra\_Shukla** 1 year, 11 months ago

Adilkhan how many questions from ET?

upvoted 1 times

🗨️ 👤 **Eltooth** 3 years, 1 month ago

Going with 4 VM's and Isolated.

upvoted 1 times

🗨️ 👤 **veenet** 3 years, 3 months ago

the resource must be allocated as Isolate

upvoted 1 times

🗨️ 👤 **yifado5713** 3 years, 6 months ago

Many questions are not from the dumps so prepare well before appearing for the exam. I passed with 720 score.

upvoted 5 times

🗨️ 👤 **sari67** 3 years, 6 months ago

I passed the exam today with score 882. All the case studies and questions were the same with these samples in examtopics, except for two or three questions and one case study related to logic app. Good luck!

upvoted 11 times

🗨️ 👤 **john4p** 3 years, 6 months ago

Nice, thanks for the info.

upvoted 1 times

🗨️ 👤 **sangmin214** 3 years, 7 months ago

4/Isolated

upvoted 1 times

🗨️ 👤 **prabhjot** 4 years, 1 month ago

yes the ans is correct

upvoted 2 times

🗨️ 👤 **mlantonis** 4 years, 1 month ago

Box 1: 4

There are four customers that use this service, and each instance of the WebJob processes data for a single customer and must run as a singleton instance. So, the number of VM should be 4.

Box 2: Isolated

Azure resources must be located in an isolated network .

In the Isolated tier, the App Service Environment defines the number of isolated workers that run your apps, and each worker is charged. In addition, there's a flat Stamp Fee for the running the App Service Environment itself. Isolated: This tier runs dedicated Azure VMs on dedicated Azure Virtual Networks. It provides network isolation on top of compute isolation to your apps. It provides the maximum scale-out capabilities.

Reference:

<https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use>

upvoted 7 times

## DRAG DROP -

You are a developer for a software as a service (SaaS) company that uses an Azure Function to process orders. The Azure Function currently runs on an Azure

Function app that is triggered by an Azure Storage queue.

You are preparing to migrate the Azure Function to Kubernetes using Kubernetes-based Event Driven Autoscaling (KEDA).

You need to configure Kubernetes Custom Resource Definitions (CRD) for the Azure Function.

Which CRDs should you configure? To answer, drag the appropriate CRD types to the correct locations. Each CRD type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

### Answer Area

CRD types	Setting	CRD type
Secret	Azure Function code	
Deployment	Polling interval	
ScaledObject	Azure Storage connection string	
TriggerAuthentication		

#### Suggested Answer:

### Answer Area

CRD types	Setting	CRD type
Secret	Azure Function code	Deployment
Deployment	Polling interval	ScaledObject
ScaledObject	Azure Storage connection string	Secret
TriggerAuthentication		

#### Box 1: Deployment -

To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

#### Box 2: ScaledObject -

With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Example of ScaledObject with polling interval

```
apiVersion: keda.k8s.io/v1alpha1
```

```
kind: ScaledObject
```

```
metadata:
```

```
name: transformer-fn
```

```
namespace: tt
```

```
labels:
```

```
deploymentName: transformer-fn
```

```
spec:
```

```
scaleTargetRef:
```

```
deploymentName: transformer-fn
```

```
pollingInterval: 5
```

```
minReplicaCount: 0
maxReplicaCount: 100
```

Box 3: Secret -

Store connection strings in Kubernetes Secrets.

Example: to create the Secret in our demo Namespace:

```
# create the k8s demo namespace
```

```
kubectl create namespace tt
```

```
# grab connection string from Azure Service Bus
```

```
KEDA_SCALER_CONNECTION_STRING=$(az servicebus queue authorization-rule keys list \
```

```
-g $RG_NAME \
```

```
--namespace-name $SBN_NAME \
```

```
--queue-name inbound \
```

```
-n keda-scaler \
```

```
--query "primaryConnectionString" \
```

```
-o tsv)
```

```
# create the kubernetes secret
```

```
kubectl create secret generic tt-keda-auth \
```

```
--from-literal KedaScaler=$KEDA_SCALER_CONNECTION_STRING \
```

```
--namespace tt
```

Reference:

<https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/>

 **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: Deployment

To deploy Azure Functions to Kubernetes use the func kubernetes deploy command has several attributes that directly control how our app scales, once it is deployed to Kubernetes.

Box 2: ScaledObject

With --polling-interval, we can control the interval used by KEDA to check Azure Service Bus Queue for messages.

Box 3: Secret


Store connection strings in Kubernetes Secrets.

upvoted 78 times

 **neerajupadhyay82** Highly Voted 4 years, 5 months ago

kubernetes is marked as out of scope but still there are few questions coming in the exam from this section

upvoted 23 times

 **Imobiletun** Most Recent 6 days, 4 hours ago

For Box 3 : In Kubernetes, a Secret is where the actual sensitive data is stored, such as an Azure Storage connection string.

But KEDA does not directly use Secret CRDs when scaling.

Instead, KEDA uses TriggerAuthentication to reference the secret securely and bind it to a specific trigger (e.g., Azure Storage Queue, Azure Service Bus, Kafka, etc.).

upvoted 1 times

 **8ac3742** 11 months ago

while the pollingInterval itself doesn't directly scale the application, it influences how often KEDA evaluates the need to scale based on the defined triggers.

upvoted 1 times

 **AhmedAbdelAziz** 1 year, 5 months ago


I think kubernetes out of the exam

upvoted 5 times

 **spahno** 1 year, 3 months ago

I hope so...

upvoted 1 times

 **jakobste** 1 year, 10 months ago



I hope this is not on the exam anymore.

upvoted 2 times

🗳️ 👤 **wrak** 2 years, 4 months ago

out-of-the-scope for the AZ-204 Exam; please remove from the Question Bank

upvoted 6 times

🗳️ 👤 **Esward** 2 years, 5 months ago

Given answers are correct!!

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-kubernetes-keda>

upvoted 1 times

🗳️ 👤 **[Removed]** 2 years, 9 months ago

This question is not AKS, but Functions being run in Kubernetes. So, technically in scope

upvoted 4 times

🗳️ 👤 **nuiowmokka** 2 years, 10 months ago

Is Kubernetes part of AZ-204, I did not study Kubernetes?

upvoted 4 times

🗳️ 👤 **VACAS** 2 years, 10 months ago

I dont know but it doesn't appers here <https://docs.microsoft.com/en-us/certifications/exams/az-204>

upvoted 4 times

🗳️ 👤 **dookebroyle** 3 years ago

Is there any Microsoft documentation on this? I can't find it.

upvoted 1 times

🗳️ 👤 **Idiw02** 3 years ago

<https://keda.sh/>

upvoted 2 times

🗳️ 👤 **Dev666** 3 years, 5 months ago

Box 1: Deployment -

Box 2: ScaledObject -

Box 3: Secret

upvoted 2 times

🗳️ 👤 **finnishr** 3 years, 10 months ago

1. Azure Function code - Deployment

-To deploy Azure Functions to Kubernetes use the func kubernetes deploy command

2. Polling interval - ScaledObject

- This is the interval to check each trigger on. By default KEDA will check each trigger source on every ScaleObject every 30 seconds.

3. Azure Storage connection string - Secret

- Store connection string in Kubernetes secret

Source for ScaledObject: <https://keda.sh/docs/1.4/concepts/scaling-deployments/>

upvoted 6 times

🗳️ 👤 **glam** 4 years, 1 month ago

Box 1: Deployment -

Box 2: ScaledObject -

Box 3: Secret -

upvoted 3 times

🗳️ 👤 **Kiranvnnv** 4 years, 7 months ago

<https://www.thinktecture.com/en/kubernetes/serverless-workloads-with-keda/serverless-workloads-with-keda/> is the right link

upvoted 5 times

🗳️ 👤 **Leandromellor** 4 years, 7 months ago



is this question in the exam?

upvoted 4 times

🗳️ 👤 **notjon** 4 years, 7 months ago

yes it is, got it yesterday when taking the 204

upvoted 15 times

  **dirtygooback** 4 years, 7 months ago

Yes. There are several Kubernetes questions. I found this odd since all the study material from MS states they use the container registry.  
upvoted 8 times

## HOTSPOT -

You are creating a CLI script that creates an Azure web app and related services in Azure App Service. The web app uses the following variables:

Variable name	Value
\$gitrepo	https://github.com/Contos/webapp
\$webappname	Webapp1103

You need to automatically deploy code from GitHub to the newly created web app.

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
az group create --location westeurope --name myResourceGroup
```

▼ --name \$webappname --resource-group myResourceGroup --sku FREE

- az webapp
- az appservice plan create
- az webapp deployment
- az group delete

▼ --name \$webappname --resource-group myResourceGroup

- az webapp create
- az appservice plan create
- az webapp deployment
- az group delete

▼

```
--repo-url $gitrepo --branch master --manual-integration
git clone $gitrepo
--plan $webappname
```

▼ source config --name \$webappname

- az webapp
- az appservice plan create
- az webapp deployment
- az group delete

```
--resource-group myResourceGroup
```

▼

```
--repo-url $gitrepo --branch master --manual-integration
git clone $gitrepo
--plan $webappname
```

Suggested Answer:

### Answer Area

```
az group create --location westeurope --name myResourceGroup
```

```
--name $webappname --resource-group myResourceGroup --sku FREE
```

az webapp
az appservice plan create
az webapp deployment
az group delete

```
--name $webappname --resource-group myResourceGroup
```

az webapp create
az appservice plan create
az webapp deployment
az group delete

```
--repo-url $gitrepo --branch master --manual-integration
```

git clone \$gitrepo
--plan \$webappname

```
source config --name $webappname
```

az webapp
az appservice plan create
az webapp deployment
az group delete

```
--resource-group myResourceGroup
```

--repo-url \$gitrepo --branch master --manual-integration
git clone \$gitrepo
--plan \$webappname

Box 1: az appservice plan create

The azure group creates command successfully returns JSON result. Now we can use resource group to create a azure app service plan

Box 2: az webapp create -

Create a new web app..

Box 3: --plan \$webappname -

..with the serviceplan we created in step 1.

Box 4: az webapp deployment -

Continuous Delivery with GitHub. Example:

```
az webapp deployment source config --name firstsamplewebsite1 --resource-group websites--repo-url $gitrepo --branch master --git-token $token
```

Box 5: --repo-url \$gitrepo --branch master --manual-integration

Reference:

<https://medium.com/@satish1v/devops-your-way-to-azure-web-apps-with-azure-cli-206ed4b3e9b1>

 **kemtin**  4 years, 1 month ago

Given answer is correct, got this on my test yesterday

upvoted 78 times

 **130nk3r5** 12 months ago

Got this today.

Went with answer here.



Score 927

upvoted 5 times

 **TonyMel** 1 year, 9 months ago

correct, in 2023Mar24, score: 904/1000

upvoted 11 times

 **Wanni**  4 years ago

```
#!/bin/bash
```

```
# Replace the following URL with a public GitHub repo URL
gitrepo=https://github.com/Azure-Samples/php-docs-hello-world
webappname=mywebapp$RANDOM
```

```
# Create a resource group.
az group create --location westeurope --name myResourceGroup
```

```
# Create an App Service plan in `FREE` tier.
az appservice plan create --name $webappname --resource-group myResourceGroup --sku FREE
```

```
# Create a web app.
az webapp create --name $webappname --resource-group myResourceGroup --plan $webappname
```

```
# Deploy code from a public GitHub repository.
az webapp deployment source config --name $webappname --resource-group myResourceGroup \
--repo-url $gitrepo --branch master --manual-integration
```



```
# Copy the result of the following command into a browser to see the web app.
echo http://$webappname.azurewebsites.net
upvoted 32 times
```

  **ozai16**  9 months, 2 weeks ago



3/12/2024 Got this on my exam, I went with the given answer  
upvoted 3 times

  **DaveOla** 9 months, 2 weeks ago



How many questions were from here?  
upvoted 1 times

  **Vukile** 1 year, 1 month ago

On exam 9 Nov 2023, went with given answer, socre 865. Case Study: Farmers and Distributors  
upvoted 2 times

  **Avi28** 1 year, 1 month ago

can you pls share what percentage of questions were from the ET dumps?  
upvoted 1 times

  **ENGs** 1 year, 2 months ago

On my exam 2023-10 before the Update of the Exam  
upvoted 2 times

  **SangeethaRamasamy** 1 year, 3 months ago

Got this in 14/09/2023 exam. This exam topic questions are helpful to pass the exam. But we need to understand the scenarios clearly so that we can able to answer the question if its twisted slightly.  
upvoted 3 times

  **70PineApple** 1 year, 10 months ago

Got this in exam today..20/02/23  
upvoted 2 times

  **Priya0703** 1 year, 10 months ago



Got this question on 20-02-2023 exam.  
upvoted 1 times

  **carlosghosn** 2 years, 1 month ago



Got this in the exam today ! Nov 25, 2022  
upvoted 5 times

  **AbdulMannan** 2 years, 3 months ago



Got this question on 30-Sep-2022 exam.  
Answer is correct. Passed with 870 score.  
upvoted 6 times

-   **coffecold** 2 years, 3 months ago



<https://learn.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-github>

upvoted 5 times
-   **serpevi** 2 years, 3 months ago



Got this in 09/22 , went with the most voted answers, score 927.

upvoted 4 times
-   **ischava780225** 2 years, 4 months ago



This is coming in the exam,

upvoted 1 times
-   **Baskman** 2 years, 9 months ago



Got this in the exam 03/22

upvoted 2 times
-   **Alasmindas** 2 years, 9 months ago

Got this in 03/22 , went with the given answers, the answers are right

upvoted 2 times
-   **oescm** 2 years, 10 months ago

Got this one 02/2022. Went with given answer.

upvoted 5 times
-   **Mev4953** 2 years, 11 months ago

Got this in the exam 01/22

upvoted 5 times

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure

Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Trigger the photo processing from Blob storage events.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

You need to catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.

Note: Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow.

Events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener.

However, the processing must start in less than one minute.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

*Community vote distribution*



**YahyaSonmez** Highly Voted 4 years, 8 months ago

Hi All,

The answer (B) is correct. Because, the trick is in the "less than one minute" detail.

You can read about "...10-minute delay in processing new blobs.." in "3-Minimizing latency" description.

Microsoft says: ".....Use Event Grid instead of the Blob storage trigger for the following scenarios:"

1-Blob-only storage accounts: Blob-only storage accounts are supported for blob input and output bindings but not for blob triggers.

2-High-scale: High scale can be loosely defined as containers that have more than 100,000 blobs in them or storage accounts that have more than 100 blob updates per second.

3-Minimizing latency: If your function app is on the Consumption plan, there can be up to a ##10-minute delay in processing new blobs## if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled. You can also use an Event Grid trigger with your Blob storage account. For an example, see the Event Grid tutorial.

REFENCE: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp#event-grid-trigger>

I wish you a good day.

upvoted 169 times

**Zidimirite** 4 years, 2 months ago

You are right, since this is a MS exam they always want to show off the variations that you can do with their platform and them specifying v2 is definitely a tell... But to get around the latency you could also turn on "Always on", which looking at that alone you could say A is right. This is not a great question.

upvoted 5 times

**Juanlu** 4 years, 8 months ago

Totally agree, YahyaSonmez !

upvoted 1 times

🗨️ 👤 **krz1** 4 years, 1 month ago

But "Blob storage events are pushed using Azure Event Grid to subscribers such as Azure Functions,"  
upvoted 3 times

🗨️ 👤 **Bartimaeus** 3 years, 4 months ago

I believe your understanding is incorrect - the answer says:  
"Trigger the photo processing from Blob storage events."  
Therefore there's no mention about actually using "Blob storage trigger".

Also, in the description of the trigger there's this note:  
"Polling works as a hybrid between inspecting logs and running periodic container scans."  
So the trigger actually uses the blob analytics logs, not blob storage events.  
It's in fact Event Grid that uses the blob storage events.

Reference: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp#polling>  
upvoted 3 times

🗨️ 👤 **[Removed]** Highly Voted 👍 5 years ago

Answer is not at all clear from the explanation. It seems like answer should be yes instead of no.  
upvoted 26 times

🗨️ 👤 **xRiot007** 5 years ago

The answer is correct. "No". You need an Azure Function for processing.  
upvoted 9 times

🗨️ 👤 **Camios** 4 years, 11 months ago

Except the question says what triggers the processing, not what does the processing. Question/answer is poorly written.  
upvoted 35 times

🗨️ 👤 **angelsrp** 4 years, 11 months ago

you can capture an event without processing the data related to the event. Given answer is correct  
upvoted 3 times

🗨️ 👤 **adelaTM** Most Recent 🔍 3 months ago

Selected Answer: A

Yes, the solution meets the goal.  
upvoted 1 times

🗨️ 👤 **SuktiSubhra** 4 months, 1 week ago

Selected Answer: A

Azure Blob Storage events can trigger an Azure Function or an Event Grid subscription whenever a new photo is uploaded.  
This ensures that the photo processing starts in less than one minute, meeting the requirement.  
Blob storage events integrate with Azure Event Grid, which can then trigger processing logic in Azure Functions, Logic Apps, or Event Hubs.  
upvoted 1 times

🗨️ 👤 **xidrane** 1 year, 4 months ago

Yes, triggering the photo processing from Blob storage events does meet the goal.

Azure Blob Storage can emit events that can be captured by Azure Event Grid. These events can then trigger various processes or workflows. In this scenario, when a user uploads a photo to Blob Storage, an event is generated. This event can be subscribed to by an Azure Function or a Logic App, which in turn can start the process of creating a mobile-friendly version of the image.

This approach is efficient and can easily meet the requirement of starting the image processing in less than one minute after the photo is uploaded. Azure Event Grid is known for its low latency in delivering events, typically in the order of seconds, which aligns well with the specified requirement.  
upvoted 3 times

🗨️ 👤 **SSR999** 1 year, 5 months ago

I think it should be 'B', as per Microsoft for blob events 'There is no service level agreement around the time it takes for a message to arrive. It's not uncommon for messages to arrive anywhere from 30 minutes to two hours.'  
<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview#practices-for-consuming-events>  
upvoted 2 times

🗨️ 👤 **bgbgvfvf** 1 year, 6 months ago

I think B is the right answer



upvoted 1 times

🗨️ 👤 **LJNZ2049** 1 year, 6 months ago

**Selected Answer: A**

I think the given answer is correct.

In this question, there is no mention of how to process the images after they are triggered, whether to use the Function app or not. It's only about the trigger itself. So the Blob storage events should meet the goal.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 11 months ago

Correct answer is A

upvoted 1 times

🗨️ 👤 **DonH** 2 years ago

Just as information I just had this question on my AZ204 exam - 16-jun-2023.

I barely made it (with only 767 points) so I can't inform anyone if this answer is correct or not, just stating that this is an actual exam question.

upvoted 4 times

🗨️ 👤 **Bababar** 2 years, 1 month ago

Correct Answer : B

Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..

Azure Storage events allow applications to react to events.

Common Blob storage event scenarios include image or video

processing, search indexing, or any file- oriented workflow.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

upvoted 1 times

🗨️ 👤 **CODE\_STS** 2 years, 4 months ago

Got this in the exam today! Feb 28, 2023

upvoted 3 times

🗨️ 👤 **fkaracan** 2 years, 4 months ago

I think "##must start## in less than one minute" is the keyword here. Thus answer is highly B.

upvoted 1 times

🗨️ 👤 **ABHI\_SHEK\_31** 2 years, 4 months ago

Yes, the solution of triggering the photo processing from Blob storage events meets the goal of starting the photo processing in less than one minute after a photo is uploaded. Azure Storage Blob storage provides event notifications for blob creation, which can be used to trigger the photo processing. This allows for the processing to start as soon as a photo is uploaded, thereby meeting the requirement of starting the process in less than one minute.

upvoted 1 times

🗨️ 👤 **Rosu123** 2 years, 5 months ago

The question states that you have developed a service. So why wait for the Blob storage to trigger when the service intercepts the event requested before even reaching the blob. So the answer is B.

upvoted 2 times

🗨️ 👤 **carlosghosn** 2 years, 7 months ago

Got this in the exam today ! Nov 25, 2022

upvoted 2 times

🗨️ 👤 **hubekpeter** 2 years, 7 months ago

**Selected Answer: A**

Please stop with that 10 minutes non sense. They mentioned event based blob storage which have low latency <https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=in-process%2Cextensionv5&pivots=programming-language-python> so the answer I A.

upvoted 4 times

🗨️ 👤 **Neckbread** 1 year ago

Whilst there is a low latency Blob storage option they do state that: "It is recommended that you use the event-based implementation as it has lower latency than the other"

upvoted 1 times

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 develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Does the solution meet the goal?

A. No

B. Yes

**Suggested Answer: A**

Specify custom warm-up.

Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot. Here's a sample web.config fragment.

```
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostname="[app hostname]" />
<add initializationPage="/Home/About" hostname="[app hostname]" />
</applicationInitialization>
</system.webServer>
```

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

Community vote distribution

B (77%)

A (23%)

 **GMartinez** Highly Voted 3 years, 1 month ago

**Selected Answer: B**

This should work:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up>

upvoted 27 times

 **TonyMel** 2 years, 3 months ago

correct, in 2023Mar24, score: 904/1000

upvoted 4 times

 **d79** 2 years, 2 months ago

The answer might be right, but the question is wrong. Auto-swap needs to be set on the source slot, not the target slot (production in this question)

upvoted 3 times

 **Lauri\_n** Highly Voted 2 years, 6 months ago

**Selected Answer: A**

I think the correct answer really might be A (No). It's not because of the last part about the warm-up being wrong. The incorrect part is before that: "You enable auto swap on the Production deployment slot.". The auto swap is enabled on the source slot, not the destination slot.

<https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots#configure-auto-swap>

upvoted 8 times

 **AAcademy** 2 years, 5 months ago

The target slot (production) is swapped with the source slot (staging). You have to warm up the staging slot with the new code and then switch it with the production slot. This way, the production slot will use the previously warmed staging slot.

upvoted 1 times

🗨️ 👤 **8ac3742** Most Recent 11 months ago

In order to ensure the required resources are available, you can prepare one pre-swap script to check if all required resources are available, you can call that pre-swap script in "az webapp config" or in CI/CD .yaml file.

upvoted 1 times

🗨️ 👤 **pjatx** 1 year, 3 months ago

Selected Answer: A

You enable auto-swap on source slot not the target slot. The enabling auto-swap on Production simply won't work with deployments to Testing slot.

upvoted 1 times

🗨️ 👤 **Isoldhe** 1 year, 4 months ago

Selected Answer: A

Copilot AI:

The solution you proposed does not meet the goal. Updating the web.config file to include the applicationInitialization configuration element only ensures that the app is warmed up before the swap, but it does not guarantee that the scripts run and the resources are available<sup>1</sup>. To achieve that, you need to use the slot swap operation API<sup>2</sup> and specify a swap operation id<sup>3</sup> that can be used to monitor the progress and status of the swap. Alternatively, you can also disable auto swap, update the app with a method named statuscheck to run the scripts, re-enable auto swap and deploy the app to the Production slot<sup>4</sup>.

upvoted 1 times

🗨️ 👤 **manopeydaikon** 1 year, 5 months ago

If answer B is correct then the question lacks some info about "warm up" part and if that part exists on the actual question on exam go with B otherwise it is correct with A.

upvoted 2 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 6 months ago

I got this question on my exam, 2023Dec, go with I remember was the most voted answer. Score 902, most of the questions were here, slightly different on wording because the Azure Ad <-> Entra Id change. Case was City Power & Light. Good luck!

upvoted 5 times

🗨️ 👤 **shekhar11** 1 year, 9 months ago

On my exam in sept23 2023

upvoted 1 times

🗨️ 👤 **Aneeqa** 1 year, 9 months ago

what is the correct ans?

upvoted 1 times

🗨️ 👤 **Tarajee** 1 year, 9 months ago

On my exam in 2023 sept10

upvoted 2 times

🗨️ 👤 **longnguyendh** 1 year, 11 months ago

It should be yes. Keyword is applicationInitialization

upvoted 2 times

🗨️ 👤 **deathRac3** 2 years, 1 month ago

Selected Answer: B

sequence was wrong

upvoted 1 times

🗨️ 👤 **p2006** 2 years, 2 months ago

Selected Answer: B

<https://learn.microsoft.com/en-us/answers/questions/980267/what-do-we-need-to-warmup-before-swapping-apps-ser>

upvoted 1 times

🗨️ 👤 **Priya0703** 2 years, 4 months ago

Got this question series on 20-02-2023 exam.

upvoted 3 times


🗨️ 👤 **Poops** 2 years, 6 months ago

Selected Answer: B

Should be B (Yes), that's the way to implement warmup according to MS:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up>

upvoted 3 times

  **gds03** 2 years, 10 months ago

You need to ensure that scripts run and resources are available before a swap operation occurs.

App Initialization has nothing to do with have sure that resources are available.

upvoted 2 times

  **rohitpatil113** 2 years, 10 months ago

**Selected Answer: A**

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#Auto-Swap>

upvoted 2 times

  **d\_vaibhav** 2 years, 10 months ago

Correct answer must be B (applicationinitialization tag is way of implementing custom warm-up)

upvoted 2 times

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 develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Enable auto swap for the Testing slot. Deploy the app to the Testing slot.

Does the solution meet the goal?

A. No

B. Yes

#### Suggested Answer: B

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot.


Here's a sample web.config fragment.

```
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostName="[app hostname]" />
<add initializationPage="/Home/About" hostName="[app hostname]" />
</applicationInitialization>
</system.webServer>
```

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

Community vote distribution



**GMartinez** Highly Voted 2 years, 7 months ago

Would say "No"

upvoted 16 times

**TonyMel** 1 year, 9 months ago

correct, in 2023Mar24, score: 904/1000

upvoted 9 times

**Mousavi** 2 years, 6 months ago

I think the given answer is correct

<https://www.pragimtech.com/blog/azure/azure-deployment-slots-auto-swap/>

upvoted 3 times

**Spartaco** Highly Voted 2 years, 7 months ago

Selected Answer: A

I vote A, No, because for me the solution is updating the web.config file to include the applicationInitialization configuration element.

upvoted 8 times

**pjatz** Most Recent 9 months, 1 week ago

Selected Answer: B

Sure Yes is the answer. Enabling auto-swap for Testing slot will do the work

"When auto swap is enabled from a slot into production, every time you push your code changes to that slot, App Service automatically swaps the app into production after it's warmed up in the source slot."

upvoted 1 times

**jessica\_\_** 9 months, 4 weeks ago

The answer is "yes"

Deploying an app to a slot first and swapping it into production makes sure that all instances of the slot are warmed up before being swapped into production. This eliminates downtime when you deploy your app. The traffic redirection is seamless, and no requests are dropped because of swap operations. You can automate this entire workflow by configuring auto swap when pre-swap validation isn't needed.

<https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots?tabs=portal>

upvoted 4 times

🗳️ 👤 **Isoldhe** 10 months, 1 week ago

Selected Answer: A

Copilot AI:

The solution you proposed does not meet the goal. Enabling auto swap for the Testing slot will only swap the Testing and Production slots when a new app is deployed to the Testing slot<sup>1</sup>. It will not ensure that the scripts run and the resources are available before the swap operation occurs. To achieve that, you need to use the slot swap operation API<sup>2</sup> and specify a swap operation id<sup>3</sup> that can be used to monitor the progress and status of the swap. Alternatively, you can also disable auto swap, update the app with a method named statuscheck to run the scripts, re-enable auto swap and deploy the app to the Production slot<sup>4</sup>.

upvoted 3 times

🗳️ 👤 **AbhishekSinha** 11 months, 2 weeks ago

Answer is No, only enabling the auto swap will not solve the problem, Question is asking for warm up action like ensure resource before the swap for that we need to provide some configuration in web.config. This configuration part is not in the solution statement

upvoted 3 times

🗳️ 👤 **Ciupaz** 11 months, 2 weeks ago

Selected Answer: A

No. Also the explanation says so.

upvoted 1 times

🗳️ 👤 **manopeydakon** 12 months ago

B, Enabling auto swap for the Testing slot and deploying the app to that slot ensures that the scripts run and resources are available in the Testing environment before swapping to Production. This allows you to validate the changes in a staging environment before they are applied to the Production slot.

upvoted 1 times

🗳️ 👤 **dddddd111** 1 year, 1 month ago

Selected Answer: B

I think the answer is "Yes". I found a very useful information about this question in this link <https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots?tabs=portal#troubleshoot-swaps>.

Search this part "Deploying your application to a nonproduction slot has the following benefits"

I know this is a series of question, but I think both this option and applicationInitialization are "Yes".

upvoted 1 times

🗳️ 👤 **Tarajee** 1 year, 3 months ago

Got this on 2023sept10

upvoted 2 times

🗳️ 👤 **Vladimir\_Gajinov** 1 year, 3 months ago

Selected Answer: A

My vote goes to 'No'.

In the "How to configure auto swap" section, step 3 "Navigate to the \*\*source slot\*\*, in our case staging slot." gives us an answer.

<https://www.pragimtech.com/blog/azure/azure-deployment-slots-auto-swap/>

upvoted 1 times

🗳️ 👤 **basquiat** 1 year, 4 months ago

All the respect for you guys commenting here, but just left: "The correct answer is X" without saying why or providing documentation kinda sucks.

upvoted 8 times

🗳️ 👤 **MakMao** 1 year, 4 months ago

Given answer is correct:

'Before you configure auto swap for the production slot, consider testing auto swap on a nonproduction target slot'

<https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots?tabs=portal>

upvoted 5 times

🗨️ 👤 **dddddd111** 1 year, 1 month ago

Seen that note on the link and updated year 2023. So confirmed that the answer is YES.

upvoted 1 times

🗨️ 👤 **VijayNarayanan** 1 year, 4 months ago

Thanks MakMao. Confusion is cleared.

upvoted 2 times

🗨️ 👤 **JH81** 1 year, 6 months ago

**Selected Answer: A**

Got this on 6/28/2023 and passed with 850. Went with No.

upvoted 3 times

🗨️ 👤 **red0400** 1 year, 10 months ago

B. Yes

Enabling auto swap for the Testing slot and deploying the app to the Testing slot will ensure that the app is fully tested and resources are available before the swap occurs to the Production slot. This way, any issues that arise during testing can be addressed before the app is made live in Production.

upvoted 7 times

🗨️ 👤 **Mike5** 1 year, 10 months ago

**Selected Answer: A**

To meet the goal of ensuring that scripts run and resources are available before a swap operation occurs, you should update the web.config file to include the application Initialization configuration element, as described in my previous answer.

upvoted 1 times

🗨️ 👤 **mdg3501** 2 years ago

got this on 2022-12-7

upvoted 3 times

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 develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Disable auto swap. Update the app with a method named statuscheck to run the scripts. Re-enable auto swap and deploy the app to the Production slot.

Does the solution meet the goal?

- A. No
- B. Yes

**Suggested Answer: B**

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot.


Here's a sample web.config fragment.

```
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostname="[app hostname]" />
<add initializationPage="/Home/About" hostname="[app hostname]" />
</applicationInitialization>
</system.webServer>
```

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

Community vote distribution



**robin1337** Highly Voted 2 years, 5 months ago

It depends.... has the app an WEBSITE\_SWAP\_WARMUP\_PING\_PATH set to /statuscheck??

Seriously, I hold multiple certifications, tricky hands-on ones but this certification is a joke. Microsoft, who is designing those questions?!? It's not just this one, there are so many bad questions.

upvoted 25 times

**[Removed]** 2 years, 3 months ago

Same experience here. I have many AWS certifications, OCP-11 certification, etc. This is really a joke of a certification. I do not even know what Microsoft expects and more importantly the kind of people ready to live with this kind of questions from Microsoft and trying to remember it.

upvoted 11 times

**robin1337** 2 years, 5 months ago

Ok, my bad autoswap has to be configured on the Staging Slot. So "You enable auto swap on the Production deployment slot." is not what you would do. Answer is A (NO).

upvoted 9 times

**GMartinez** Highly Voted 2 years, 7 months ago

**Selected Answer: A**

Instead, use applicationInitialization

upvoted 22 times

**TonyMel** 1 year, 9 months ago

correct, in 2023Mar24, score: 904/1000

upvoted 5 times

**Isoldhe** Most Recent 10 months, 1 week ago



Selected Answer: B

Copilot AI:

The solution you proposed does meet the goal. Disabling auto swap, updating the app with a method named statuscheck to run the scripts, re-enabling auto swap and deploying the app to the Production slot is one of the possible ways to ensure that scripts run and resources are available before a swap operation occurs. This method uses the applicationInitialization configuration element in the web.config file to specify the statuscheck method as a custom initialization action. The statuscheck method will run the scripts and check the availability of the resources before the app is swapped.

upvoted 3 times

🗳️ 👤 **manopeyda** 12 months ago

How can it say:

Correct Answer: B

then, "Instead" update the web.config file to include the applicationInitialization configuration element..

Sounds very wrong!

upvoted 3 times

🗳️ 👤 **Tarajee** 1 year, 3 months ago

Got this on 2023sept

upvoted 1 times

🗳️ 👤 **GuruPratap** 1 year, 9 months ago

Yes, the solution should meet the goal of ensuring that scripts run and resources are available before a swap operation occurs. By disabling auto swap, the deployment to the Production slot will not occur until it is manually initiated. Then, the app can be updated with a method named statuscheck to run the necessary scripts and ensure that resources are available. Once the app has been updated, auto swap can be re-enabled and the updated app can be deployed to the Production slot. This approach allows for greater control over when the swap occurs and ensures that the necessary steps are taken before the swap happens.

upvoted 5 times

🗳️ 👤 **Saronyo** 1 year, 10 months ago

Selected B , Application initialisation can be set up to do the statuscheck . auto swap will deploy it to production after warm up

upvoted 3 times

🗳️ 👤 **red0400** 1 year, 10 months ago

Exactly

upvoted 1 times

🗳️ 👤 **kociuba** 2 years ago

got in on my exam 30-12-2022

score: 818

upvoted 2 times

🗳️ 👤 **Congphan** 1 year, 12 months ago

what is your choice?

upvoted 2 times

🗳️ 👤 **mdg3501** 2 years ago

got this on 2022-12-7

upvoted 1 times

🗳️ 👤 **micro9000** 2 years, 1 month ago

The answer is A, you still need to set the WEBSITE\_WARMUP\_PATH in app settings

WEBSITE\_WARMUP\_PATH: A relative path on the site that should be pinged whenever the site restarts (not only during slot swaps). Example values include /statuscheck or the root path, /.

<https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots#Auto-Swap>

upvoted 2 times

🗳️ 👤 **[Removed]** 2 years, 3 months ago

Deploy the app directly to the Production slot will fail the auto swap. But that is a detail in the question that will be missed by a person who did some reading on the auto swap because he is thinking of app initialization option vs custom\_warmup with those properties. Do not be evil, Microsoft.

upvoted 2 times

🗳️ 👤 **N9** 2 years, 4 months ago

Selected Answer: A

web.config file with applicationInitialization is correct answer. So A is the correct answer.

upvoted 2 times

  **Dani\_ac7** 2 years, 6 months ago

**Selected Answer: B**

applicationInitialization

upvoted 2 times

  **Lucky\_me** 2 years, 6 months ago

**Selected Answer: B**

although the provided explanation in the answer section, is not related to the correct option, however, the solution provided in the question would actually work, and that's the correct "Manual Way" of making sure the deployment slot is warmed up and running before swap, so the solution would work and correct.

upvoted 3 times

  **Satish\_Babu** 2 years, 6 months ago

**Selected Answer: A**

No,

No,

because correct solution is updating the web.config file to include applicationInitialization configuration element for running initializing resources.

upvoted 1 times

  **cemepo9819** 2 years, 5 months ago

multiple solutions can be correct.

upvoted 4 times



  **omega11** 2 years, 6 months ago

I say A, "No".

The reason why setting up auto swap is to warm up the application to prevent downtime, so deploying to the production slot directly is killing the purpose.

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up>

upvoted 3 times

  **Eltooth** 2 years, 6 months ago

**Selected Answer: A**

A seems to be correct answer.

upvoted 2 times

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure

Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Convert the Azure Storage account to a BlockBlobStorage storage account.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: B

Not necessary to convert the account, instead move photo processing to an Azure Function triggered from the blob upload..


Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file- oriented workflow.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

Community vote distribution



**oleks** Highly Voted 4 years, 4 months ago

B is correct

upvoted 32 times

**MiraA** 3 years, 2 months ago

The answer is NO - the Even Grid integration should be used.

<https://docs.microsoft.com/en-us/azure/event-grid/resize-images-on-storage-blob-upload-event>

upvoted 10 times

**Justing\_Gao** Highly Voted 4 years, 5 months ago

A BlockBlobStorage account is a specialized storage account in the premium performance tier for storing unstructured object data as block blobs or append blobs. Compared with general-purpose v2 and BlobStorage accounts, BlockBlobStorage accounts provide low, consistent latency and higher transaction rates.

upvoted 17 times

**syj123** 4 years, 4 months ago

so, Do you mean that BlockBlobStorage account is better but not necessary?

upvoted 2 times

**pjatx** Most Recent 9 months, 1 week ago

Selected Answer: B

Azure storage of any type won't help here. You need something that will trigger processing. Answer NO.

upvoted 1 times

**bgbgvfvf** 1 year ago

I think B is the right answer

upvoted 1 times

**kvtborad** 1 year, 4 months ago

I have this question on 6th August. it was passed with 904. chose the NO.

upvoted 4 times

- 🗄️ 👤 **[Removed]** 1 year, 8 months ago  
Got this in exam today (5 April 2023)  
upvoted 2 times
- 🗄️ 👤 **CODE\_STS** 1 year, 10 months ago  
Got this in the exam today! Feb 28, 2023  
upvoted 1 times
- 🗄️ 👤 **ABHI\_SHEK\_31** 1 year, 10 months ago  
No, converting the Azure Storage account to a BlockBlobStorage storage account does not meet the requirement of starting the photo processing in less than one minute. The conversion of the storage account type would not have any impact on the time it takes to start the photo processing. A different solution such as triggering the photo processing from Blob storage events or using a queue-based solution may be more appropriate to meet the requirement.  
upvoted 4 times
- 🗄️ 👤 **kociuba** 2 years ago  
got it on my exam 30-12-2022  
score: 818  
upvoted 1 times
- 🗄️ 👤 **SoftSol** 2 years ago  
**Selected Answer: B**  
Got this in the exam dec 13, 2022  
upvoted 2 times
- 🗄️ 👤 **carlosghosn** 2 years, 1 month ago  
Got this in the exam today ! Nov 25, 2022  
upvoted 1 times
- 🗄️ 👤 **AZAdam22** 2 years, 5 months ago  
**Selected Answer: B**  
No - The account type does not affect how the photo processing code is triggered.  
upvoted 6 times
- 🗄️ 👤 **dookebroyle** 2 years, 6 months ago  
**Selected Answer: A**  
No, because you can't convert this account. I think the question is intentionally misleading, wanting you to be distracted by the 1 minute start to process time. Note that the requirements state the processing needs to START in less than 1 minute, not finish in a minute.  
upvoted 1 times
- 🗄️ 👤 **Eltooth** 2 years, 6 months ago  
**Selected Answer: B**  
B is correct answer.  
upvoted 1 times
- 🗄️ 👤 **Goyo** 2 years, 10 months ago  
**Selected Answer: B**  
The correct answer is NO  
upvoted 1 times
- 🗄️ 👤 **dmeld** 3 years, 1 month ago  
Answer is No. Blob storage trigger cannot handle the requirement less than one minutes. You should use the Event Grid trigger. Hence, it is not useful change the account storage tier from v2 to block blob account.  
upvoted 2 times
- 🗄️ 👤 **Chked** 3 years, 5 months ago  
Answer - NO  
"You can't convert an existing standard performance storage account to a block blob storage account with premium performance. To migrate to a premium performance storage account, you must create a premium block blob account, and migrate the data to the new account."  
<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-performance-tiers>  
upvoted 10 times

## HOTSPOT -

You are developing an Azure Web App. You configure TLS mutual authentication for the web app.

You need to validate the client certificate in the web app. To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Property	Value
Client certificate location	<div>▼</div> <div>HTTP request header</div> <div>Client cookie</div> <div>HTTP message body</div> <div>URL query string</div>
Encoding type	<div>▼</div> <div>HTML</div> <div>URL</div> <div>Unicode</div> <div>Base64</div>

### Answer Area

Suggested Answer:

Property	Value
Client certificate location	<div>▼</div> <div>HTTP request header</div> <div>Client cookie</div> <div>HTTP message body</div> <div>URL query string</div>
Encoding type	<div>▼</div> <div>HTML</div> <div>URL</div> <div>Unicode</div> <div>Base64</div>

Accessing the client certificate from App Service.

If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the `HttpRequest.ClientCertificate` property. For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth>

**mlantonis** Highly Voted 3 years, 7 months ago

Box 1: HTTP request header

If you are using ASP.NET and configure your app to use client certificate authentication, the certificate will be available through the `HttpRequest.ClientCertificate` property.

Box 2: Base64

For other application stacks, the client cert will be available in your app through a base64 encoded value in the "X-ARR-ClientCert" request header. Your application can create a certificate from this value and then use it for authentication and authorization purposes in your application.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth>

upvoted 67 times

🗲️ 👤 **27close** Highly Voted 🍌 4 years, 1 month ago

With client certificates enabled, App Service injects an X-ARR-ClientCert request header with the client certificate.

upvoted 16 times

🗲️ 👤 **neelkanths** Most Recent 🔔 8 months ago

Got it on 20 April 2024...Marks > 950...answer is correct....

Box 1: HTTP request header

Box 2: Base64

upvoted 6 times

🗲️ 👤 **Nihilist11** 11 months, 2 weeks ago

Jan-21-2024 - This was in my exam, went with selected answer Score 740

Contosso case study [Couldnt find here]

upvoted 3 times

🗲️ 👤 **bgbgvfvf** 1 year ago

Answer is correct.Base64

upvoted 1 times

🗲️ 👤 **NightshadeRC** 1 year, 5 months ago

Had this question today: 2023-07-26

upvoted 4 times

🗲️ 👤 **[Removed]** 1 year, 8 months ago

Got this in exam today (5 April 2023)

upvoted 3 times

🗲️ 👤 **ucskips** 1 year, 9 months ago

why is it base64?

upvoted 3 times

🗲️ 👤 **yyandrakk** 1 year ago

"or other application stacks (Node.js, PHP, etc.), the client cert is available in your app through a base64 encoded value in the X-ARR-ClientCert request header."

upvoted 1 times

🗲️ 👤 **CODE\_STS** 1 year, 10 months ago

Got this in the exam today! Feb 28, 2023

upvoted 1 times

🗲️ 👤 **uffuchsi** 1 year, 10 months ago

Received this in my exam today (22/02/2023). Selected HTTP request header, and Base64. Score 927.

upvoted 4 times

🗲️ 👤 **Jyo221** 1 year, 10 months ago

Are the questions coming from dump

upvoted 1 times

🗲️ 👤 **Esward** 1 year, 10 months ago

It is there in 13 Feb 2023 exam

upvoted 2 times

🗲️ 👤 **zb1234** 2 years ago

Got this Dec 28 2022

upvoted 4 times

🗲️ 👤 **nvtienanh** 2 years ago

Got this question on 2-Dec-2022 exam.

Answer is correct. Passed with 857 score.

upvoted 6 times

🗲️ 👤 **ms\_master** 2 years, 2 months ago

Got this question on 30-Oct-2022 exam.

Answer is correct. Passed with 875 score



upvoted 2 times

  **AbdulMannan** 2 years, 3 months ago

Got this question on 30-Sep-2022 exam.



Answer is correct. Passed with 870 score.

upvoted 2 times

  **serpevi** 2 years, 3 months ago

Got this in 09/22 , went with the most voted answers, score 927.

upvoted 2 times

  **N9** 2 years, 4 months ago

Box 1: HTTP request header

Box 2: Base64

upvoted 2 times

## DRAG DROP -

You are developing a Docker/Go using Azure App Service Web App for Containers. You plan to run the container in an App Service on Linux. You identify a

Docker container image to use.

None of your current resource groups reside in a location that supports Linux. You must minimize the number of resource groups required.

You need to create the application and perform an initial deployment.

Which three Azure CLI commands should you use to develop the solution? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

**Azure CLI Commands****Answer Area**





**Azure CLI Commands****Answer Area**









Suggested Answer:



You can host native Linux applications in the cloud by using Azure Web Apps. To create a Web App for Containers, you must run Azure CLI commands that create a group, then a service plan, and finally the web app itself.

Step 1: az group create -

In the Cloud Shell, create a resource group with the az group create command.

Step 2: az appservice plan create

In the Cloud Shell, create an App Service plan in the resource group with the az appservice plan create command.

Step 3: az webapp create -

In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command. Don't forget to replace with a unique app name, and <docker-ID> with your Docker ID.

Reference:

<https://docs.microsoft.com/mt-mt/azure/app-service/containers/quickstart-docker-go?view=sql-server-ver15>

**Secure01** 3 years, 7 months ago

Within the same resource group, you can't mix Windows and Linux apps in the same region.



<https://docs.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux>



"None of your current resource groups reside in a location that supports Linux"

So you have to create new resource group. Answer is correct.

upvoted 71 times

  **rgullini** 2 years, 9 months ago

This explanation is correct and it is applicable BEFORE January 21, 2021. Now the article says:

"Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this"

upvoted 27 times

  **kags** 2 years, 8 months ago

Agreed. Just created both linux and windows, core3.1 and .net 5 respectively in same resource group.

upvoted 4 times

  **Netspu** 2 years, 4 months ago

Many are missing the point here, it says "None of your current resource groups reside in a location that supports Linux". So you need to create a group in an area that does support Linux. Doesn't matter if you mix them or not, there is currently nothing that supports Linux, so something new is required.

For me the answer is correct.

upvoted 33 times

  **bubbleGumStroganoff** 2 years, 3 months ago

The resources inside of a resource group can be in a region that differs from the region of the resource group they reside in though no?

upvoted 6 times

  **rohitpatil113** 1 year, 10 months ago

Limitations

Note

Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this.

upvoted 15 times

  **Edwardenis** 3 years, 3 months ago

I don't agree with you.

1. the link you posted doesn't say anything about restriction on mixing App service plan in the same region.

2. I just created 2 app service plan in the same RG. One app service plan on windows in the region of the RG East US and the other app service plan in Linux in Central US

upvoted 1 times

  **azurelearner666** 3 years ago

It does if you read it... see under limitations, it says "You can't mix Windows and Linux apps in the same App Service plan."

upvoted 6 times

  **dookebroyle** 2 years ago

It also states your current resource group locations do not support Linux. So even though you can in your situation, the question specifies that that is not an option

upvoted 4 times

  **JulienYork** 3 years, 5 months ago

<https://docs.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux>

upvoted 2 times

  **Ide**  3 years, 4 months ago

<https://docs.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux>

Historically, you can't mix Windows and Linux apps in the same resource group. However, all resource groups created on or after January 21, 2021 do support this scenario. For resource groups created before January 21, 2021, the ability to add mixed platform deployments will be rolled out across Azure regions (including National cloud regions) soon.

upvoted 20 times

  **azurelearner666** 3 years ago

This should be updated on the exam to reflect reality... so we have to respond wrong now to pass... LOL

upvoted 6 times

🗳️ 👤 **Edwardenis** 3 years, 3 months ago

Ohhh Thanks for the notice, thats why I could create both app service plans.

upvoted 2 times

🗳️ 👤 **pandugadu009** Most Recent 9 months, 1 week ago

the given answer is correct. the az group update command can not update the location. so you need to create a new group.

upvoted 5 times

🗳️ 👤 **juanckar** 11 months, 4 weeks ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 14 times

🗳️ 👤 **surprise0011** 1 year, 2 months ago

you can't update the resource group, you will get an error.

So it has to be: az group create

az group update --name acrrg --set location=centralus

(InvalidResourceGroupLocation) Invalid resource group location 'centralus'. The Resource group already exists in location 'westeurope'.

Code: InvalidResourceGroupLocation

Message: Invalid resource group location 'centralus'. The Resource group already exists in location 'westeurope'.

upvoted 4 times

🗳️ 👤 **ucskips** 1 year, 3 months ago

Once the region of a resource group is set it cannot be updated. I read this before but I don't where

upvoted 3 times

🗳️ 👤 **Esward** 1 year, 4 months ago

It was there in 13 Feb 2023 exam

upvoted 1 times

🗳️ 👤 **obeesubz5** 1 year, 5 months ago

All you need here is to create an app service plan and the web app.

Secondly, no one said the existing rg contain Windows vm's.

Thirdly, it does not say that the existing rg's don't support linux, it says "the locations" of the rg don't support linux. So here you have to bare in mind that a resource need not be in the same region as the rg. You can google, in fact try it out yourself. This supports the "Limit number of resource groups" requirement.

upvoted 3 times

🗳️ 👤 **Esward** 1 year, 5 months ago

Given answers are correct!!

upvoted 1 times

🗳️ 👤 **ghassen204** 1 year, 6 months ago

I think:

1: az group update

2: az appservice plan create

3: az webapp create

Note: Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this.

upvoted 4 times

🗳️ 👤 **josecouva** 1 year, 5 months ago

But in the queston it says "None of your current resource groups reside in a location that supports Linux", so you can't execute az group update. I

think the answer given is correct

upvoted 4 times

🗳️ 👤 **[Removed]** 1 year, 9 months ago

It says the Resource groups are not in a location that supports Linux, but you can create a resource in a location other than the resource group location. So, I do not understand why there is a need to create a new RG. But I do agree that given the options the answer is correct

upvoted 5 times

🗳️ 👤 **Enigma\_\_** 1 year, 9 months ago

Linux and Windows App Service plans can now share resource groups. This limitation has been lifted from the platform and existing resource groups have been updated to support this.

<https://learn.microsoft.com/en-us/azure/app-service/overview#app-service-on-linux>

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 10 months ago

\* The question says ResourceGroup is residing in a location that does not support Linux. But where the resource-group resides has no relation to where the Resources (App service, for example) is residing.

\* Reduce the number of resource group. The totally useless comment like that is just given to create confusion for the benefit of the certification industry

The whole discussion in the comments about Windows and Linux App service plan cannot co locate in an RG is totally unrelated. Nobody said about a Windows app service plan. Either the question is wrong or Microsoft just cannot get anything right

upvoted 2 times

🗨️ 👤 **xRiot007** 2 years ago

1. Create group (it very states that your current resources do \*not\* support Linux)
2. Create service plan
3. Create the web app

upvoted 6 times

🗨️ 👤 **silvatyrant** 2 years, 4 months ago

- 1) Linux and Windows apps can now be deployed in the same RG.
- 2) They mention doing an initial deployment as well. "az webapp up" is best suited to do this. Not "az webapp create". The "create" action will not do an initial deployment.

upvoted 3 times

🗨️ 👤 **dookebroyle** 2 years ago

I agree, the "right" answer doesn't deploy the app. Webapp up will create a new resource group and app service plan though, so the answer is the same. Your current resource group locations don't support Linux, so your only option was to move them (not supported) or create a new one.

upvoted 1 times

🗨️ 👤 **Mev4953** 2 years, 5 months ago

Got this in the exam 01/22

upvoted 3 times

🗨️ 👤 **Mev4953** 2 years, 5 months ago

No, sorry i didnt. I read wrong :)

upvoted 2 times

🗨️ 👤 **resonant** 2 years, 6 months ago

A couple of things to note in this question:

1. It says "Which three Azure CLI commands should you use to develop the solution?" which means that you HAVE to choose 3 options. Therefore, even if it looks like you don't have to create or update a resource group, that should be enough hint that at least you should choose one of them.
2. You can't change the location of a resource group once it is created. Not even with the "az group update" command.

With 1. and 2. it should be clear enough that you have to create the resource group anyways, even if you didn't pay attention to the fact that the question says none of the locations of the resource groups already created can have Linux (independently of the fact that Linux and Windows can NOW coexist in the same resource group).

upvoted 13 times

🗨️ 👤 **AmineD** 1 year, 7 months ago

I totally agree with the fact that resource group location cannot be changed after creation and since "None of your current resource groups reside in a location that supports Linux" so we have to use "az group create" instead of "az group update"

For me the answer is correct

Reference: <https://learn.microsoft.com/en-us/rest/api/resources/resource-groups/create-or-update>

upvoted 2 times

## DRAG DROP -

Fourth Coffee has an ASP.NET Core web app that runs in Docker. The app is mapped to the `www.fourthcoffee.com` domain.

Fourth Coffee is migrating this application to Azure.

You need to provision an App Service Web App to host this docker image and map the custom domain to the App Service web app.

A resource group named `FourthCoffeePublicWebResourceGroup` has been created in the WestUS region that contains an App Service Plan named `AppServiceLinuxDockerPlan`.

Which order should the CLI commands be used to develop the solution? To answer, move all of the Azure CLI commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

**Azure CLI Commands**

```
az webapp config container set
--docker-custom-image-name
$dockerHubContainerPath
--name $appName
--resource-group
fourthCoffeePublicWebResourceGroup
```

```
az webapp config hostname add
--webapp-name $appName
--resource-group
fourthCoffeePublicWebResourceGroup \
--hostname $fqdn
```

```
az webapp create
--name $appName
--plan AppServiceLinuxDockerPlan
--resource-group
fourthCoffeePublicWebResourceGroup
```

```
#!/bin/bash
appName="FourthCoffeePublicWeb$random"
location="WestUS"
dockerHubContainerPath="FourthCoffee/publicweb:v1"
fqdn="http://www.fourthcoffee.com">www.fourthcoffee.com
```

**Answer Area**

Suggested Answer:

### Azure CLI Commands

```
az webapp config container set
--docker-custom-image-name
$dockerHubContainerPath
--name $appName
--resource-group
fourthCoffeePublicWebResourceGroup
```

```
az webapp config hostname add
--webapp-name $appName
--resource-group
fourthCoffeePublicWebResourceGroup \
--hostname $fqdn
```

```
az webapp create
--name $appName
--plan AppServiceLinuxDockerPlan
--resource-group
fourthCoffeePublicWebResourceGroup
```

```
#!/bin/bash
appName="FourthCoffeePublicWeb$random"
location="WestUS"
dockerHubContainerPath="FourthCoffee/publicweb:v1"
fqdn="http://www.fourthcoffee.com">www.fourthcoffee.com
```

### Answer Area

```
#!/bin/bash
appName="FourthCoffeePublicWeb$random"
location="WestUS"
dockerHubContainerPath="FourthCoffee/publicweb:v1"
fqdn="http://www.fourthcoffee.com">www.fourthcoffee.com
```

```
az webapp create
--name $appName
--plan AppServiceLinuxDockerPlan
--resource-group
fourthCoffeePublicWebResourceGroup
```

```
az webapp config container set
--docker-custom-image-name
$dockerHubContainerPath
--name $appName
--resource-group
fourthCoffeePublicWebResourceGroup
```

```
az webapp config hostname add
--webapp-name $appName
--resource-group
fourthCoffeePublicWebResourceGroup \
--hostname $fqdn
```

Step 1: #bin/bash -

The appName is used when the webapp-name is created in step 2.

Step 2: az webapp create -

Create a web app. In the Cloud Shell, create a web app in the myAppServicePlan App Service plan with the az webapp create command.

Step 3: az webapp config container set

In Create a web app, you specified an image on Docker Hub in the az webapp create command. This is good enough for a public image. To use a private image, you need to configure your Docker account ID and password in your Azure web app.


Step 4: az webapp config hostname add

The webapp-name is used when the webapp is created in step 2.


In the Cloud Shell, follow the az webapp create command with az webapp config container set.


Reference:

<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image> <https://docs.microsoft.com/en-us/azure/app-service/tutorial-custom-container?pivot=container-linux> <https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-configure-custom-domain>

 **msuser11**  2 years, 7 months ago

1. /bin/bash
  2. az webpp create
  3. ~ config container set
  4. ~ config hostname add
- upvoted 113 times

 **surprise0011** 1 year, 8 months ago  
received 2023-04-17 went above answer, score 926  
upvoted 10 times

 **ayush66203** 1 year, 2 months ago  
were all the questions in the exam from exam topic? did you prepare from paid version of exam topic or the free one?  
upvoted 2 times

 **nmatnoh** 2 years, 6 months ago

I would agree with the propose solution

1. /bin/bash
2. az webpp create
3. ~ config container set

URL: <https://docs.microsoft.com/en-us/azure/app-service/tutorial-custom-container?pivots=container-linux>

At the bottom of tutorial, link to next tutorial

In the next tutorial, you learn how to map a custom DNS name to your app.

4. ~ config hostname add

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-custom-domain>

upvoted 5 times

  **Mousavi** 2 years, 6 months ago

I still have doubts about the order of 3 and 4 ?

upvoted 7 times

  **GMartinez**  2 years, 7 months ago

Shouldn' it be az webapp create as 2nd, az webapp config container set as 3rd and az webapp config hostname as last ?

upvoted 11 times

  **nightflyer** 2 years, 7 months ago

Yes this is correct, i saw this question in the Microsoft Learn Labs

upvoted 4 times

  **examtopicsLogin123**  4 months, 2 weeks ago

Got it in February 2025. Went with highly voted answer.

upvoted 2 times

  **abhishekgraphs** 7 months, 1 week ago

This was on my exam. Date: 23rd May 2024.

Went with given answer, got score of 925.

Out of 43 questions asked, approximately 38 were from ExamTopics questions. I had purchased contributor access to AZ-204 in ExamTopics

I am more concerned about those 5 questions not from this question bank. I think those were the ones that I probably got wrong in the test. Hence the score of 925. Else would have scored higher had all 43 questions been from this question bank

upvoted 4 times

  **neelkanths** 8 months ago

Got it on 20 April 2024...Marks > 950...answer is correct....all questions from examtopics 400 question bank

upvoted 5 times

  **onlyforheros** 9 months, 2 weeks ago

Got it in the exam 13.03.2024. Score: 910. Went with the give answer



upvoted 1 times

  **Nihilist11** 11 months, 2 weeks ago

Jan-21-2024 - This was in my exam, went with selected answer Score 740

Contosso case study [Couldnt find here]

upvoted 1 times

  **1989jay** 11 months, 3 weeks ago

This question was on exam. 09/01/2024. Passed 872

upvoted 1 times

  **arunkuml** 1 year ago

Got it in the exam 14/12/23. Went with given answer. All questions are from ExamTopics. Case study - VanArsdel, Ltd (11 questions)

upvoted 2 times

  **CarlosTheBoldest** 1 year ago

I got this question on my exam, 2023Dec, go with I remember was the most voted answer. Score 902, most of the questions were here, slightly different on wording because the Azure Ad <-> Entra Id change. Case was City Power & Light. Good luck!

upvoted 1 times

  **AndySmith** 1 year, 1 month ago

On exam 3-Nov-2023. Went with proposed answer - 932/1000.

upvoted 1 times

🗨️ 👤 **JH81** 1 year, 6 months ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.

upvoted 3 times

🗨️ 👤 **Maximuskh** 1 year, 6 months ago

I think the reason container needs to be set and then hostname is because if we configure hostname first while the app does not have docker image configuration, the end users experience downtime.

upvoted 3 times

🗨️ 👤 **aragones** 1 year, 7 months ago

Got this 2023-05-12.

FYI:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 2 times

🗨️ 👤 **Saluk\_DE** 1 year, 9 months ago

Question was in Exam 2023-03-30

upvoted 2 times

🗨️ 👤 **sarmaria** 1 year, 9 months ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 2 times

🗨️ 👤 **Esward** 1 year, 11 months ago

Given answers are correct!

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-custom-domain>

upvoted 1 times



## DRAG DROP -

You are developing a serverless Java application on Azure. You create a new Azure Key Vault to work with secrets from a new Azure Functions application.

The application must meet the following requirements:

- ⇒ Reference the Azure Key Vault without requiring any changes to the Java code.
- ⇒ Dynamically add and remove instances of the Azure Functions host based on the number of incoming application events.
- ⇒ Ensure that instances are perpetually warm to avoid any cold starts.
- ⇒ Connect to a VNet.
- ⇒ Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.

You need to grant the Azure Functions application access to the Azure Key Vault.

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

### Answer Area

Create a user-assigned managed identity for the application.

Create the Azure Functions app with a Premium plan type.

Create an access policy in Azure Key Vault for the application identity.

Create an SSL certification in Azure Key Vault for the application identity.

Create the Azure Functions app with an App Service plan type.

Create the Azure Functions app with a Consumption plan type.

Create a system-assigned managed identity for the application.



Suggested Answer:

### Actions

Create a user-assigned managed identity for the application.

Create the Azure Functions app with a Premium plan type.

Create an access policy in Azure Key Vault for the application identity.

Create an SSL certification in Azure Key Vault for the application identity.

Create the Azure Functions app with an App Service plan type.

Create the Azure Functions app with a Consumption plan type.

Create a system-assigned managed identity for the application.



### Answer Area

Create the Azure Functions app with a Consumption plan type.

Create a user-assigned managed identity for the application.

Create an access policy in Azure Key Vault for the application identity.



Step 1: Create the Azure Functions app with a Consumption plan type.

Use the Consumption plan for serverless.

Step 2: Create a system-assigned managed identity for the application.

Create a system-assigned managed identity for your application.

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Step 3: Create an access policy in Key Vault for the application identity.

Create an access policy in Key Vault for the application identity you created earlier. Enable the "Get" secret permission on this policy. Do not



configure the  
"authorized application" or applicationId settings, as this is not compatible with a managed identity.  
Reference:  
<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>

msuser11 Highly Voted 2 years, 7 months ago

1. create ~Premium plan Type (Consumption X)
  2. create system-assigned ~ (user-assigned X)
  3. create an access policy in Azure Key Vault~
- upvoted 180 times

NaSit 2 years, 6 months ago

I agree with you.

1. Premium plan (avoid any cold starts and connect to a VNet)

Overview of plans here: <https://docs.microsoft.com/th-th/azure/azure-functions/functions-scale>

2. create system-assigned => "A system-assigned identity is tied to your application and is deleted if your app is deleted."
3. create an access policy

<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?toc=%2Fazure%2Fazure-functions%2Ftoc.json&tabs=azure-cli>

upvoted 26 times

Xardas 1 year, 5 months ago

Why not App Service plan?

upvoted 1 times

macobuzi 1 year, 4 months ago

App Service plan can connect to Vnet but it won't "Ensure that instances are perpetually warm to avoid any cold starts".

upvoted 4 times

imheretolearn 1 year, 3 months ago

Wouldn't Always On setting on App Service plan work here?

upvoted 3 times

jakobste 1 year, 4 months ago

I agreee. I would go with system assigned identity unless the question talks about "multiple apps/vms". In that case you would go for user assigned if "administration has to be minimized" etc.

upvoted 1 times

fkaracan 1 year, 10 months ago

correct

upvoted 1 times

Esward 1 year, 11 months ago

Agreed

upvoted 1 times

uffuchsi Highly Voted 1 year, 10 months ago

Received this in my exam today (22/02/2023). Selected 'Created the Azure Function app with Premium plan type', 'Create a system-assigned managed identity for the application', and 'Create an access policy in Azure Key Vault for the application identity'. Score 927.

upvoted 25 times

J\_obando13 Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

raja33 2 months, 1 week ago

Got this in the exam today! Apr 19, 2025

upvoted 1 times

cb98160 9 months, 2 weeks ago

1. Premium plan (avoid cold start)
2. system-assigned managed identity (we only have 1 Azure resource that our webapp needs to access and also the managed identity has to be

removed when we delete or app.)

3. Create an access policy in Azure Key Vault.

upvoted 6 times

🗨️ 👤 **AhmedAbdelAziz** 11 months ago

Why there is a lot of incorrect answers this website show to us ?

upvoted 3 times

🗨️ 👤 **Samuel316** 1 year, 2 months ago

Answer shown in the image does not agree with what's written in the answer description. Image says user-assigned managed identity. Description says system-assigned managed identity, which would be more correct

upvoted 6 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 1 month ago

System assigned, as the user assigned continue existing after the resource is removed, but system assigned won't be valid once the resource is deleted

upvoted 2 times

🗨️ 👤 **Tarajee** 1 year, 3 months ago

Got this on 2023sept

upvoted 2 times

🗨️ 👤 **kvtborad** 1 year, 4 months ago

I have this question on 6th August. it was passed with 904. chose this highly voted:create ~Premium plan Type (Consumption X)

2. create system-assigned ~ (user-assigned X)

3. create an access policy in Azure Key Vault~

upvoted 4 times

🗨️ 👤 **tom112** 1 year, 7 months ago

system-assigned Managed Identity's life cycle:

Shared life cycle with the Azure resource that the managed identity is created with.

When the parent resource is deleted, the managed identity is deleted as well.

upvoted 1 times

🗨️ 👤 **ucskips** 1 year, 9 months ago

I agree with the answers here 1. Create function app with premium plan (premium plan allows for virtual network connection and pre-warmed starts)

2. System assigned managed identity (when the function app is deleted so is the system identity they are "tied" together

<https://devblogs.microsoft.com/devops/demystifying-service-principals-managed-identities/>) 3. create an access policy (here, the system identity can be set as the method of accessing the key vault\*\*)

upvoted 1 times

🗨️ 👤 **ucskips** 1 year, 9 months ago

I agree with the answers here

1. Create function app with premium plan (premium plan allows for virtual network connection and pre-warmed starts)

2. System assigned managed identity (when the function app is deleted so is the system identity they are "tied" together

<https://devblogs.microsoft.com/devops/demystifying-service-principals-managed-identities/>)

3. create an access policy (here, the system identity can be set as the method of accessing the key value)

upvoted 1 times

🗨️ 👤 **CODE\_STS** 1 year, 10 months ago

Got this in the exam today! Feb 28, 2023

upvoted 1 times

🗨️ 👤 **Harish86** 1 year, 10 months ago

Can you tell me, why this website is showing incorrect answers? And most importantly from where it is this answers. When it is saying about Warm tiers , how can it be Consumption plan

upvoted 4 times

🗨️ 👤 **adilkhan** 1 year, 11 months ago

consumption plan does not support VNET <https://learn.microsoft.com/en-us/azure/azure-functions/functions-networking-options?tabs=azure-cli>

upvoted 2 times

🗨️ 👤 **micro9000** 2 years, 1 month ago

The Azure Function App must be in Premium plan because we need to ensure that instances are perpetually warm to avoid any cold starts.

upvoted 2 times

 **vruirob** 2 years, 1 month ago

The following requirements are the key:

⇒ Ensure that instances are perpetually warm to avoid any cold starts.

In this link, <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#cold-start-behavior>, you can read this same phrase, so, the correct answer is Premium Plan

⇒ Authentication to the Azure Key Vault instance must be removed if the Azure Function application is deleted.

In this link, <https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types>, you can see that is the System-assigned the correct answer

⇒ And access policy in Azure Key Vault

upvoted 4 times

You develop a website. You plan to host the website in Azure. You expect the website to experience high traffic volumes after it is published. You must ensure that the website remains available and responsive while minimizing cost. You need to deploy the website. What should you do?

- A. Deploy the website to a virtual machine. Configure the virtual machine to automatically scale when the CPU load is high.
- B. Deploy the website to an App Service that uses the Shared service tier. Configure the App Service plan to automatically scale when the CPU load is high.
- C. Deploy the website to a virtual machine. Configure a Scale Set to increase the virtual machine instance count when the CPU load is high.
- D. Deploy the website to an App Service that uses the Standard service tier. Configure the App Service plan to automatically scale when the CPU load is high.

**Suggested Answer: D**

Windows Azure Web Sites (WAWS) offers 3 modes: Standard, Free, and Shared.

Standard mode carries an enterprise-grade SLA (Service Level Agreement) of 99.9% monthly, even for sites with just one instance.

Standard mode runs on dedicated instances, making it different from the other ways to buy Windows Azure Web Sites.

Incorrect Answers:

B: Shared and Free modes do not offer the scaling flexibility of Standard, and they have some important limits.

Shared mode, just as the name states, also uses shared Compute resources, and also has a CPU limit. So, while neither Free nor Shared is likely to be the best choice for your production environment due to these limits.


Community vote distribution

D (100%)

  **profesorklaus** Highly Voted 4 years ago

Answer is correct. No doubts!

upvoted 94 times

  **27close** Highly Voted 4 years, 1 month ago

answer is d- this is production and requires autoscaling when cpu is peak. the rest is constraints by limit without auto scaling

upvoted 22 times

  **onlyforheros** Most Recent 9 months, 2 weeks ago

Got it in the exam 13.03.2024. Score: 910. Went with the give answer

upvoted 4 times

  **AlbertoBT** 1 year, 1 month ago

C Is the right one.

"Standard" Tier does "not allow" Automatic scaling.

Automatic scaling requerres Premium V2 or Premium V3 tier

upvoted 1 times

  **CarlosTheBoldest** 1 year, 1 month ago

It does, check it out here:

<https://azure.microsoft.com/en-us/pricing/details/app-service/windows/#pricing>

upvoted 3 times

  **AlbertoBT** 1 year, 1 month ago

I was wrong. "Standard" Tier does "not allow" Automatic scaling BASED ON TRAFFIC

But is allow BASED ON METRICS (Like CPU suage)

upvoted 4 times

  **AlbertoBT** 1 year, 1 month ago

I was right.

Is not clear because microsoft naming

- Scale out method "AUTOMATIC" if to scale based on traffic

- Scale out method "RULES BASED" if to scale based on metrics

In fact is not possible create a AUTOMATIC scaling based on METRICS

<https://learn.microsoft.com/en-us/azure/app-service/manage-automatic-scaling?tabs=azure-portal>  
upvoted 2 times

🗨️ 👤 **AlbertoBT** 1 year, 1 month ago  
C: Scale Set. Is the right answer.

<https://learn.microsoft.com/en-us/azure/virtual-machine-scale-sets/virtual-machine-scale-sets-autoscale-overview#autoscale-rule-criteria>  
upvoted 2 times

🗨️ 👤 **YodaR10** 1 year, 1 month ago  
Why not B?  
upvoted 1 times

🗨️ 👤 **RuffBoii** 1 year, 3 months ago  
Had this on my exam today.  
upvoted 3 times

🗨️ 👤 **juanckar** 1 year, 5 months ago  
This was on the exam (July 2023). Went with proposed. Scored 917  
upvoted 3 times

🗨️ 👤 **alexein74** 1 year, 11 months ago

**Selected Answer: D**

D - To ensure that the website remains available and responsive while minimizing cost, you should deploy the website to an App Service that uses the Standard service tier and configure the App Service plan to automatically scale when the CPU load is high. This way, the website can handle high traffic volumes by automatically scaling the number of instances of the website, reducing the risk of the website becoming unavailable due to high traffic.  
upvoted 6 times

🗨️ 👤 **ReyPirata** 1 year, 4 months ago  
correct, in 08/20/2023, score: 925/1000  
upvoted 1 times

🗨️ 👤 **TonyMel** 1 year, 9 months ago  
correct, in 2023Mar24, score: 904/1000  
upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 1 month ago

**Selected Answer: D**

Answer is correct.  
upvoted 3 times

🗨️ 👤 **nhannn** 2 years, 3 months ago

**Selected Answer: D**

The answer is D, it is very obvious both C and D are correct but we are aiming for the minimizing cost, then it should be D  
upvoted 3 times

🗨️ 👤 **kyawtlearning** 2 years, 4 months ago

The answer is correct to save the cost.  
upvoted 1 times

🗨️ 👤 **jmirza** 2 years, 5 months ago

I think the scale for the Wwe Apps should be based on Network and RAM mostly.  
upvoted 1 times



🗨️ 👤 **AZAdam22** 2 years, 5 months ago

**Selected Answer: D**

D - Because app service standard tier is the cheapest tier that supports auto scale.  
upvoted 2 times

🗨️ 👤 **dookebroyle** 2 years, 6 months ago

I think this would depend on whether you want to have to manage the VMs. If you choose VMs, you are responsible for OS and updates. I would think D would be the right answer  
upvoted 1 times

  **Eltooth** 2 years, 6 months ago

**Selected Answer: D**

D is correct answer.  
upvoted 2 times

  **drgn7676** 2 years, 8 months ago

**Selected Answer: D**

Uses the Standard pricing tier  
upvoted 2 times

  **DarkyShadoW** 2 years, 8 months ago

**Selected Answer: D**

D. AppService Standar with AutoScale  
upvoted 2 times

## HOTSPOT -

A company is developing a Java web app. The web app code is hosted in a GitHub repository located at <https://github.com/Contoso/webapp>. The web app must be evaluated before it is moved to production. You must deploy the initial code release to a deployment slot named staging. You need to create the web app and deploy the code.

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
```

az  create --location centralus --name \$resourcegroupname

group  
webapp  
appservice plan  
webapp deployment slot  
webapp deployment source

az  create --name \$webappname --resource-group \$resourcegroupname --sku S3

group  
webapp  
appservice plan  
webapp deployment slot  
webapp deployment source

az  create --name \$webappname --resource-group \$resourcegroupname --plan \$webappname

group  
webapp  
appservice plan  
webapp deployment slot  
webapp deployment source

az  create --name \$webappname --resource-group \$resourcegroupname --slot staging

group  
webapp  
appservice plan  
webapp deployment slot  
webapp deployment source


az  config --name \$webappname --resource-group \$resourcegroupname \  
--slot staging --repo-url \$gitrepo --branch master --manual-integration

group  
webapp  
appservice plan  
webapp deployment slot  
webapp deployment source


Suggested Answer:

### Answer Area


```
gitrepo=https://github.com/Contoso/webapp
webappname=businesswebapp
resourcegroupname=BusinessAppResourceGroup
```

```
az  create --location centralus --name $resourcegroupname
```


group
webapp
appservice plan
webapp deployment slot
webapp deployment source

```
az  create --name $webappname --resource-group $resourcegroupname --sku S3
```


group
webapp
appservice plan
webapp deployment slot
webapp deployment source

```
az  create --name $webappname --resource-group $resourcegroupname --plan $webappname
```

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

```
az  create --name $webappname --resource-group $resourcegroupname --slot staging
```

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

```
az  config --name $webappname --resource-group $resourcegroupname \
--slot staging --repo-url $gitrepo --branch master --manual-integration
```

group
webapp
appservice plan
webapp deployment slot
webapp deployment source

Box 1: group -

# Create a resource group.

```
az group create --location westeurope --name myResourceGroup
```

Box 2: appservice plan -

# Create an App Service plan in STANDARD tier (minimum required by deployment slots). az appservice plan create --name \$webappname --resource-group myResourceGroup --sku S1

Box 3: webapp -

# Create a web app.

```
az webapp create --name $webappname --resource-group myResourceGroup \
--plan $webappname
```

Box 4: webapp deployment slot -

#Create a deployment slot with the name "staging".

```
az webapp deployment slot create --name $webappname --resource-group myResourceGroup \
--slot staging
```

Box 5: webapp deployment source -

# Deploy sample code to "staging" slot from GitHub.

```
az webapp deployment source config --name $webappname --resource-group myResourceGroup \
--slot staging --repo-url $gitrepo --branch master --manual-integration
```

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

 **sskPL**  4 years, 7 months ago

Given answer is correct.

upvoted 201 times

 **130nk3r5** 1 year, 6 months ago



Got this today.  
Went with answer here.  
Score 927  
upvoted 8 times

🗨️ 👤 **Harshad\_Mehta** Highly Voted 4 years, 7 months ago  
why do you guys have to comment if the answer is correct.  
upvoted 35 times

🗨️ 👤 **AhmedAbdelAziz** 1 year, 5 months ago  
We all here bro for the community and the discussion I got a lot of knowledge from the community regardless the QA  
upvoted 7 times

🗨️ 👤 **azurelearner666** 4 years ago  
you commented too as well, right? ;)  
upvoted 15 times

🗨️ 👤 **DonH** 2 years ago  
No, he asked a valid question. @thomas204 gave a valid answer. I also agree on that. Some answers are wrong so it's satisfying to read the comments when an answer is correct, just to be on the safe side.  
upvoted 3 times

🗨️ 👤 **thomas204** 4 years, 7 months ago  
Because some answers are wrong and commenting when correct is just comforting for others.  
upvoted 248 times

🗨️ 👤 **thebarber87** 2 years, 3 months ago  
So vote. It is annoying when a question has 80 comments on it but they're just all answer is correct.  
upvoted 1 times

🗨️ 👤 **szydlo** 10 months, 3 weeks ago  
you cannot vote on hotspot type questions  
upvoted 2 times

🗨️ 👤 **azurelearner666** 4 years ago  
Agree! love this site because of the community and the discussions. they also help to understand the reasoning!!  
upvoted 39 times

🗨️ 👤 **lasyan3** 3 years, 9 months ago  
Totally agreed, thanks to all the community members :)  
upvoted 12 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago  
Yes, many dump test websites have these questions but I still had to move to this site for the community!  
upvoted 2 times

🗨️ 👤 **Ciupaz** 1 year, 8 months ago  
Me too.  
upvoted 1 times

🗨️ 👤 **bgbgvf** Most Recent 1 year, 6 months ago  
The given solution is correct  
upvoted 2 times

🗨️ 👤 **Vukile** 1 year, 7 months ago  
On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors  
upvoted 2 times

🗨️ 👤 **Ciupaz** 1 year, 7 months ago  
Is this a new Case Study?  
upvoted 1 times

🗨️ 👤 **NPE\_** 1 year, 7 months ago  
Got this on the exam Nov. 07 2023  
upvoted 1 times

🗨️ 👤 **nlk0** 1 year, 10 months ago

Got this today 26/08/2023, went with the given. 880  
Case study was contoso -> 9 questions  
+ 35 question, two series.  
3 of them were not from examtopics, all others were  
upvoted 2 times

🗲️ 👤 **MysticalSam** 2 years ago

This question was in today's exam at 10-June-2023  
upvoted 4 times

🗲️ 👤 **Videira** 2 years, 4 months ago

On my exam 2023-02-25  
upvoted 4 times

🗲️ 👤 **70PineApple** 2 years, 4 months ago

Got this in exam today..20/02/23  
score: 817  
upvoted 4 times

🗲️ 👤 **saravanasanthosh** 2 years, 6 months ago

Got this on 12/30/2022 exam. The given solution is correct  
upvoted 2 times

🗲️ 👤 **kociuba** 2 years, 6 months ago

got it on my exam 30-12-2022  
score: 818  
upvoted 3 times

🗲️ 👤 **Suwani** 2 years, 6 months ago

Given answer is correct  
upvoted 1 times

🗲️ 👤 **mdg3501** 2 years, 6 months ago

got this on 2022-12-7  
upvoted 1 times

🗲️ 👤 **TheExamMaster2020** 2 years, 7 months ago

Did my exam on 15th November 2022. This question was on it.  
upvoted 2 times

🗲️ 👤 **rol204** 2 years, 8 months ago

Box 1: group  
# Create a resource group.

Box 2: appservice plan  
# Create an App Service plan in STANDARD tier (minimum required by deployment slots)

Box 3: webapp  
# Create a web app

Box 4: webapp deployment slot  
#Create a deployment slot with the name "staging"

Box 5: webapp deployment source  
# Deploy sample code to "staging" slot from GitHub  
upvoted 4 times

🗲️ 👤 **Anitay** 2 years, 9 months ago

Answers are correct...  
upvoted 1 times

🗲️ 👤 **Elsheimy** 3 years, 1 month ago

On exam, 05/22  
upvoted 2 times

## HOTSPOT -

You have a web service that is used to pay for food deliveries. The web service uses Azure Cosmos DB as the data store.

You plan to add a new feature that allows users to set a tip amount. The new feature requires that a property named tip on the document in Cosmos DB must be present and contain a numeric value.

There are many existing websites and mobile apps that use the web service that will not be updated to set the tip property for some time.

How should you complete the trigger?

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

```
function ensureTip() {  
  var r = 

|                             |
|-----------------------------|
| __value();                  |
| __readDocument('item');     |
| getContext().getRequest();  |
| getContext().getResponse(); |

  
  var i = r.getBody();  


|                                             |
|---------------------------------------------|
| if (!("tip" in i)) {                        |
| if (request.getValue("tip") === null) {     |
| if (isNaN(i["tip"])    i["tip"] === null) { |
| if (typeof __.pluck("tip") === 'number') {  |

  
    i["tip"] = 0;  
  }  


|                        |
|------------------------|
| r.setBody(i);          |
| r.setValue(i);         |
| __.upsertDocument(i);  |
| __.replaceDocument(i); |

  
}
```

## Answer Area

```
function ensureTip() {  
  var r = 

|                             |
|-----------------------------|
| __value();                  |
| __readDocument('item');     |
| getContext().getRequest();  |
| getContext().getResponse(); |

  
  var i = r.getBody();  
  


|                                             |
|---------------------------------------------|
| if (!("tip" in i)) {                        |
| if (request.getValue("tip") === null) {     |
| if (isNaN(i["tip"])    i["tip"] === null) { |
| if (typeof __.pluck("tip") === 'number') {  |

  
    i["tip"] = 0;  
  }  
  


|                        |
|------------------------|
| r.setBody(i);          |
| r.setValue(i);         |
| __.upsertDocument(i);  |
| __.replaceDocument(i); |

  
}
```

Suggested Answer:

 **dol** Highly Voted 4 years, 1 month ago

Right answer in second drop down is the first one (... "tip" in i...)

Similar example can be found on <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs>  
upvoted 173 times

 **Juanlu** 4 years, 1 month ago

After develop a simple Javascript test, the correct option is (... "tip" in i). In cases where "tip" is not defined, sentence "if(isNaN(i)["tip"]) || i["tip"] === null)", return "undefined", so it doesn't work properly !

Here this simple test:

```
var i = { tip: 5, name: "just a test"}
```

```
if ('tip' in i){  
  console.log("YES");  
}  
else {  
  console.log("NO");  
}
```

```
console.log("Option 2");
```

```
if(isNaN(i)["tip"]) || i["tip"] === null){ // Undefined when "tip: N" not in array !  
  console.log("NO");  
}else {  
  console.log("YES");  
}
```

upvoted 11 times

 **nychollas** 3 years, 6 months ago

because the syntax in the proposed option is NaN(i)["tip"]... it should have been NaN(i["tip"]). In this case, this option might be taken into account as a possible correct answer. This is my opinion.

upvoted 3 times

 **Satheesh88** 4 years ago

question says tip value should be numeric. What if is gave var i = { tip: 'junk', name: "just a test"}

Result:

"YES"

"Option 2"

"NO"

upvoted 6 times

  **thomastrain** 4 years ago

Sadheesh88 has a valid point. Although the question states "there will be many apps not updating their value for some time", a good solution should always validate user input. If the value of "tip" is null or "junk", the first option will fail.

Option 3 has a typo and is missing a closing parenthesis. If we make a small change to option 3, so it reads:

```
if(isNaN(i["tip"]) || i["tip"] === null){
```

then this becomes the best answer. Check for yourself with this fiddle:

<https://jsfiddle.net/d4x3fota/2/>

So I think it comes down to if the answers were copied correctly.

upvoted 22 times

  **DonH** 1 year, 6 months ago



thanks for the Fiddle.

upvoted 2 times

  **zoan** 4 years ago

I tested the isNaN in and I get 'true' when 'tip' is undefined, null or "a" . But I use isNaN(i['tip']) and not isNaN(i)['tip']. For 'tip' in i I get only if tip is defined the i object regardless of the type of the value.

upvoted 1 times

  **Tom87** 3 years, 10 months ago

isNaN(null) is false. That's why i["tip"] === null is needed too.

upvoted 1 times

  **MikeAWS** 1 year, 2 months ago

Correct! Second dropdown is: if(!("tip" in i)) {

answer can also be found here: <https://www.whizlabs.com/blog/microsoft-azure-az-204-exam-questions//>

upvoted 2 times

  **jakobste** 1 year, 4 months ago

Yes. This question is on Wizlabs also with a variation, and the code there is (... "tip" in i...).

upvoted 1 times

  **sujitwarrier11** 3 years, 6 months ago

where is the validation that the tip value is a number? so the current answer is right. isNaN checks if the value is not a number. if true it sets value to 0.

upvoted 8 times

  **altafpatel1984** 3 years, 1 month ago

But they suggest isNaN(i), not isNaN(i["tip"]) which makes no sense as syntax is incorrect.

upvoted 6 times

  **MiraA** 3 years, 2 months ago

My opinion to NaN()...

The "old" clients will send no "tip" value at all => set "tip" to 0 to match the new expectations.

The "new" clients will send some "tip" value => pass "tip" value to a web service as it is; the web service has possibility to detect NaN() itself and to reject the request as invalid with proper HTTP status code (and log such invalid request to detect malformed clients). The client knows something went wrong and that his request wasn't processed.

I consider 0 to be a default value of "tip" which will be used for "old" clients only.

Correcting wrong (NaN) values in "tip" to 0 at this moment is bad idea as it hides a problem with the client's communication.

A reference to "a numeric value" requirement in the assignment is related to CosmosDB primarily I think.



upvoted 1 times

  **Skyrocket** Highly Voted 3 years, 11 months ago

Admin, request you to update answers to avoid any further confusion.

1. getRequest
2. (!"tip" in i)
3. setBody

upvoted 101 times

  **Forhallf** 3 years, 10 months ago

Skyrocket is correct.

```
function validateToDolItemTimestamp() {
  var context = getContext();
  var request = context.getRequest();

  // item to be created in the current operation
  var itemToCreate = request.getBody();

  // validate properties
  if (!("timestamp" in itemToCreate)) {
    var ts = new Date();
    itemToCreate["timestamp"] = ts.getTime();
  }

  // update the item that will be created
  request.setBody(itemToCreate);
}
```

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs>


upvoted 9 times

  **btrump** 3 years, 10 months ago

This is not correct. The prompt states, "must be present and contain a numeric value."



if i["tip"] == "foo" your passes but fails to meet the requirement.

upvoted 12 times

  **jasifu3** 2 years, 10 months ago

your argument is valid, but there is no better alternative. isNaN(i)["tip"] is evaluated as false["tip"] -> undefined every time.

upvoted 2 times

  **Basu525** 3 years, 11 months ago

this is the correct answer. admins please rectify it.

upvoted 4 times

  **sujitwarrier11** 3 years, 6 months ago

where is the validation that the tip value is a number? so the current answer is right. isNaN checks if the value is not a number. if true it sets value to 0.

upvoted 4 times

  **Poops** 2 years ago

isNaN(i) checks that i is NaN, it will always evaluate to false. That whole snippet is bad, it only compiles because JavaScript has a very forgiving syntax

upvoted 1 times

  **sujitwarrier11** 3 years, 6 months ago

where is the validation that the tip value is a number? so the current answer is right. isNaN checks if the value is not a number. if true it sets value to 0.

upvoted 8 times

  **HernanP** Most Recent 4 months, 1 week ago

I'll go for the following:

```
function ensureTip() {
  var r = getContext().getResponse(); // ✓ Correct selection
```

```
var i = r.getBody();
```

```
if (!("tip" in i)) { // ✓ Correct selection  
i["tip"] = 0;  
}
```

```
r.setBody(i); // ✓ Correct selection  
}
```

upvoted 1 times

🗳️ 👤 **Jak007** 11 months ago

Seems like the right answer for the second drop down should be the first alternative (... "tip" in i...)! First and second are correct. Thanks to dol for giving good link and direction.

Similar example can be found on <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-write-stored-procedures-triggers-udfs#pre-triggers>

upvoted 2 times

🗳️ 👤 **dddddd111** 1 year, 1 month ago

Check out this link with almost same sample code. <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-write-stored-procedures-triggers-udfs?tabs=javascript>

```
function validateToDoltemTimestamp() {  
var context = getContext();  
var request = context.getRequest();
```

```
// item to be created in the current operation  
var itemToCreate = request.getBody();
```

```
// validate properties  
if (!("timestamp" in itemToCreate)) {  
var ts = new Date();  
itemToCreate["timestamp"] = ts.getTime();  
}
```

```
// update the item that will be created  
request.setBody(itemToCreate);  
}
```

upvoted 3 times

🗳️ 👤 **ENGs** 1 year, 2 months ago

On my exam 2023-10 before the Update of the Exam

upvoted 2 times

🗳️ 👤 **jakobste** 1 year, 4 months ago

This was on the exam in August 2023

upvoted 4 times

🗳️ 👤 **Sleepers** 1 year, 6 months ago

Got this on 06/27/2023 exam

upvoted 4 times

🗳️ 👤 **Devendra\_Shukla** 1 year, 5 months ago

which option you select for second dropdown?

upvoted 1 times

🗳️ 👤 **DonH** 1 year, 6 months ago

Just for information: I just had this question on my AZ204 exam - 16-jun-2023.

I barely made it (with only 767 points) so I can't inform anyone if this answer is correct or not, just stating that this is an actual exam question.

upvoted 7 times

🗳️ 👤 **dirtuebus** 1 year, 7 months ago

The right answer in the second dropdown, even though it has a little syntax error, has to be by exclusion:

```
if(isNaN(i)["tip"] || i["tip"] === null) {
```

The other options don't check if "tip" is a number, at all.

Anyway, even if there is a syntax error, I don't believe the exam is supposed to be a test of our ability to compile the code in our minds :)

upvoted 5 times

🗨️ **D58** 1 year, 7 months ago

The goal is to set the tip property on the document in Cosmos DB. 3. question's answer should be `__replaceDocument(i)` and not `r.Body(i)` since the document should be replaced (updated) and not only returned to the user (`r.Body(i)`). See link with examples here: <https://learn.microsoft.com/en-us/training/modules/work-with-cosmos-db/4-cosmos-db-stored-procedures>

upvoted 1 times

🗨️ **D58** 1 year, 7 months ago

Oh never mind: Pretrigger can be used:

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-write-stored-procedures-triggers-udfs?tabs=javascript#pre-triggers>

upvoted 2 times

🗨️ **fkaracan** 1 year, 10 months ago

Question says "tip" should be numeric. Thus `isNaN(tip)` is correct

upvoted 2 times

🗨️ **dimsok** 1 year, 10 months ago

The `isNaN(i)["tip"]` makes no sense as a syntax, `isNaN` returns true or false.

From the line `i["tip"] = 0`; we can assume `i` is an array and thus ("`tip`" in `i`) is correct

upvoted 2 times

🗨️ **saravanasanthosh** 2 years ago

Got this on 12/30/2022 exam

upvoted 2 times

🗨️ **carlosghosn** 2 years, 1 month ago

Got this in the exam today ! Nov 25, 2022

upvoted 2 times

🗨️ **[Removed]** 2 years, 3 months ago

Microsoft expect us to be polygots knowing Javascript, .Net, Powershell and Bash. Or maybe this exam is to prove that the person is full time busy doing hello-world deployments using different options in Azure. A decent developer uses Azure not try to remember only Azure. This is pointless.

The only thing to know is that if other webapps are not yet ready to use that new variable (tip) then it should be set in the trigger to be 0 if not provided (validations are not strictly required because it will fail if the tip is given wrongly as a string and assuming there are no business requirement to set boundaries)

The code is not required to check this knowledge. A decent developer knows the art of googling-copy-and-paste.

upvoted 5 times

🗨️ **macobuzi** 1 year, 4 months ago

I understand your struggle, but Microsoft also doesn't expect us to study their questions bank to pass their exam. Just like a job interview, many interview questions might not relate to your actual work at all. Still, you need to pass it.

upvoted 1 times

🗨️ **r3verse** 2 years, 1 month ago

Hey man i keep seeing you post negative about every question in this exam. I do understand your points, but it doesn't add any value to the people trying to learn the questions. It actually decreaseas the value to add to it. Your feedback about Microsoft in general can be posted at lots of places, but i don't think this is the best place. Please try to contribute in a way that helps people get the exam!

upvoted 11 times

🗨️ **NombreFalso** 1 year, 10 months ago

I think he's pretty fun for comedic relief

upvoted 5 times

🗨️ **[Removed]** 2 years, 3 months ago

I'm a developer and I get a headache looking at that "code". Microsoft should stop thinking these kind of code is the only way to test a "developer".

Just look at all these discussions about this answer and what knowledge is it checking? (Memorize what the Microsoft guy wrote in some document as example?)

upvoted 6 times



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 develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob. The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Use the Durable Function async pattern to process the blob data.

Does the solution meet the goal?

- A. Yes
- B. No

**Suggested Answer: B**

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.


Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

Community vote distribution



**sasisang** Highly Voted 4 years, 7 months ago

if someone knows the exact answer or any reference , then only comment. It is so confusing  
upvoted 92 times

**john4p** 3 years, 6 months ago

A. Yes

"230 seconds is the maximum amount of time[...] For longer processing times, consider using the DURABLE FUNCTIONS ASYNC PATTERN[...]"  
<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>  
upvoted 56 times

**Netspud** 3 years, 4 months ago

See question 17, which would suggest the correct answer for this question is NO. The issues is the calling app is timing out not the function (i.e. the function is not replying as completed to the calling app in time, so use a queue, queue replies as completed in time, and work is deferred). At least after reading the two that is how I understand it. Very confusing question.  
upvoted 4 times

**xRiot007** 3 years ago

these types of questions can have multiple correct responses, and sometimes no correct responses at all. In this case, Durable Functions and using http triggers that can defer to a queue for long time processing are good answers.  
upvoted 3 times

**Magis** 2 years, 7 months ago

I agree with xRiot007. This question have two correct answers as MS saying that you can use Durable Function OR Service Bus in this scenario. So are correct.

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

upvoted 1 times

🗨️ 👤 **rgullini** 2 years, 2 months ago

Somehow you need to investigate and understand the answers, not just memorize without understanding.

upvoted 7 times

🗨️ 👤 **montoya8201** 1 year, 1 month ago

It's true

upvoted 1 times

🗨️ 👤 **surprise0011** 2 years, 2 months ago

received 2023-04-17 went with most voted answer, score 926

upvoted 4 times

🗨️ 👤 **Awry** Highly Voted 4 years, 8 months ago

Shouldn't this be TRUE?

<https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-overview?tabs=csharp#async-http>

upvoted 58 times

🗨️ 👤 **sieiro** 3 years, 7 months ago

I thing that maybe durable functions exceed de timeout of other functions, but if the functions fails or interrump don't cover the "The app must process the blob data".

upvoted 1 times

🗨️ 👤 **noros** 3 years, 4 months ago

True, I think the explanation mentions this

<https://docs.microsoft.com/en-us/azure/architecture/patterns/async-request-reply#example>

which looks exactly like the Async Pattern of Durable Function

upvoted 1 times

🗨️ 👤 **Secure01** 4 years, 7 months ago

Seems you are right...

Durable Functions provides built-in support for this pattern(Pattern #3: Async HTTP APIs), simplifying or even removing the code you need to write to interact with long-running function executions.

upvoted 8 times

🗨️ 👤 **WillPassExam** 4 years, 2 months ago

Agree :)

Durable function is one of the suggested two options for long processing tasks (the other is to defer the actual work (e.g. push to service bus queue + queue triggered function solution)

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

upvoted 28 times

🗨️ 👤 **fiercelfreedom** Most Recent 8 months ago

A function app cannot be triggered by an output binding. The wording is confusing.

upvoted 1 times

🗨️ 👤 **HariB1992** 1 year, 2 months ago

The solution provided does not meet the goal. The Durable Function async pattern is not necessary for this scenario. Instead, you can address the timeout issue by adjusting the execution timeout settings for your Azure Function app

upvoted 1 times

🗨️ 👤 **Isoldhe** 1 year, 4 months ago

Selected Answer: A

Copilot AI:

Yes, the solution meets the goal. Using the Durable Function async pattern to process the blob data can help you avoid the HTTP timeout issue and handle long-running operations in a serverless environment. Durable Functions let you write stateful functions that can orchestrate other functions and entities using the Azure Functions programming model. You can use the async HTTP API pattern to implement a polling mechanism that allows the client to check the status and the result of the operation.

upvoted 3 times

🗨️ 👤 **FeriAZ** 1 year, 4 months ago

**Selected Answer: A**

Using the Durable Function async pattern allows the function app to orchestrate the processing of blob data in a way that does not tie up the initial HTTP triggered function. Instead, the HTTP trigger can initiate the process and respond immediately, while the actual processing happens asynchronously. This method effectively circumvents the timeout issue because the durable orchestration can run for much longer than the standard function timeout period.

upvoted 1 times

🗨️ 👤 **AhmedAbdelAziz** 1 year, 5 months ago

Correct Answer No : As the app time out you have to use service bus queue and trigger function app when item added to the queue and for the better performance you can event grid and let the function trigger based on the service bus trigger events from the event grid

upvoted 1 times

🗨️ 👤 **Jak007** 1 year, 5 months ago

**Selected Answer: A**

A is true as per the documentation here: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

However the explanation in the answer provided by ExamTopics is also correct it seems.

upvoted 1 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

Using the Durable Function async pattern can help overcome the timeout limitation. Durable Functions provide a way to write stateful functions in a serverless environment, and they support long-running workflows. By breaking down the processing into smaller, manageable steps and leveraging the async pattern, you can ensure that the function continues processing the blob data without timing out. So I would say Yes.

upvoted 1 times

🗨️ 👤 **Vladimir\_Gajinov** 1 year, 9 months ago

**Selected Answer: A**

The answer is YES.

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

upvoted 2 times

🗨️ 👤 **Michael19842** 1 year, 11 months ago

This site is not helping me at all... Every third answer is wrong

upvoted 5 times

🗨️ 👤 **sarmaria** 2 years, 3 months ago

Got this question in the exam on 16/03/2023. Chosen yes. Scored 846. Make sure to prepare for case studies. I got city and lights case study.

upvoted 3 times

🗨️ 👤 **daothanhien** 2 years, 4 months ago

**Selected Answer: B**

No.

Durable changes timeout settings. It can't resolve the problem.

Use queues to response immediately

upvoted 1 times

🗨️ 👤 **jimmyml** 2 years, 4 months ago

Yes, using the Durable Function async pattern can solve the problem of the app timing out and ensure that the blob data is processed. The Durable Function allows long-running functions to be executed in a more scalable and reliable way, by splitting the functions into smaller, asynchronous tasks that can be managed by the framework. This can help prevent timeouts by allowing the function to continue processing in the background, while also providing additional benefits such as state management and error handling.

upvoted 1 times

🗨️ 👤 **dimsok** 2 years, 4 months ago



I wouldn't use durable functions, they are build to solve stateful problems, and in any case If the calling app is timing out, then durable functions will do nothing.

upvoted 3 times

🗨️ 👤 **Esward** 2 years, 5 months ago



Yes, Is the correct answer

upvoted 1 times

  **Esward** 2 years, 5 months ago

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

upvoted 1 times

  **narenazure** 2 years, 5 months ago

**Selected Answer: A**

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

upvoted 2 times

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 develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob. The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: A

Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

Community vote distribution



**msdevanms** Highly Voted 4 years, 1 month ago

Answer is correct

upvoted 66 times

**piotrek1993** 2 years, 9 months ago

What about sending files via Service Bus ? I guess this is not a good idea especially when maximum message size is 256 kb.

upvoted 3 times

**hubekpeter** 2 years, 1 month ago

You are not going to send the file thru message buss, only reference to a blob, 256kB is plenty enough for that purpose.

upvoted 5 times

**tom112** 1 year, 7 months ago

But it's a message, not an event. There are chances that a message would contain a blob in their body, max size 256k for standard and 100MB for premium.

upvoted 1 times

**vruizrob** 2 years, 3 months ago

This link <https://docs.microsoft.com/en-us/azure/azure-functions/performance-reliability#avoid-long-running-functions> confirms that the correct answer is A.

upvoted 7 times

**surprise0011** 1 year, 8 months ago

good finding. now it is clear.

IMO it not seems optimal but it is certainly a solution

upvoted 1 times

**noros** 2 years, 10 months ago

Yes, the note section here confirms that

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

upvoted 2 times

**henry1985** 2 years, 5 months ago

Async request-reply pattern <https://docs.microsoft.com/en-us/azure/architecture/patterns/async-request-reply#example>  
upvoted 2 times

🗳️ 👤 **Secure01** Highly Voted 🏆 4 years, 1 month ago

Answer is NO. The best solution is durable functions

Durable Functions provides built-in support for this pattern(Pattern #3: Async HTTP APIs), simplifying or even removing the code you need to write to interact with long-running function executions.

Durable Functions provides built-in support for this pattern(Pattern #3: Async HTTP APIs), simplifying or even removing the code you need to write to interact with long-running function executions.

<https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-overview?tabs=csharp#async-http>

upvoted 27 times

🗳️ 👤 **Figa** 4 years, 1 month ago

Why should this Answer be false?

Can't be both correct? They are not asking for the optimal solution.

upvoted 21 times

🗳️ 👤 **Cornholioz** 4 years, 1 month ago

You may be right. There are questions where two answers are true. I think both Durable Timers and this solution are both true.

upvoted 10 times

🗳️ 👤 **Cornholioz** 4 years, 1 month ago

But then again, can you really chunk the Blob data into multiples just so the refactored smaller function sets can process in parallel? What if the Blob data is just one huge chunk of a single transaction that has to be completed in shot even if takes long. That makes the Durable Functions the better/right/only answer :(

upvoted 1 times

🗳️ 👤 **tevivi8222** 3 years, 9 months ago

It was timing out because you were doing a long task during HTTP request, which can't live forever. If you're saving a huge blob and it's not being done during the lifecycle of that request, the time doesn't really matter.

upvoted 1 times

🗳️ 👤 **rgullini** 1 year, 8 months ago

Question is not asking for the BEST solution. Is asking if the provided solution will solve the issue or not, despite it is more expensive or not.

upvoted 6 times

🗳️ 👤 **xRiot007** 2 years, 6 months ago

You are asked if the solution is possible, which it is, along with using Durable Functions

upvoted 2 times

🗳️ 👤 **AOE** 3 years, 4 months ago

Both are correct: service bus queue or durable function async pattern

upvoted 21 times

🗳️ 👤 **SummerWarrior** 2 years, 11 months ago

If a service bus queue trigger is used, wouldn't the function app still be timed out? The problem seems to be the processing time of the blob data.

upvoted 5 times

🗳️ 👤 **melior** Most Recent 🕒 1 week, 5 days ago

Selected Answer: A

Yes, the proposed architecture (using an Azure Service Bus queue for asynchronous processing) does meet the goal of preventing timeouts while ensuring blob data is processed reliably.

upvoted 1 times

🗳️ 👤 **HariB1992** 8 months, 1 week ago

Yes, the solution meets the goal. By passing the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and returning an immediate HTTP success response, you can address the timeout issue and ensure that the blob data is processed without timing out.

upvoted 1 times

🗳️ 👤 **blpiek21** 9 months ago

Selected Answer: A

Correct

upvoted 1 times

🗨️ 👤 **Isoldhe** 10 months, 1 week ago

**Selected Answer: A**

Copilot AI:

Yes, the solution meets the goal. Passing the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and returning an immediate HTTP success response can help you avoid the HTTP timeout issue and handle long-running operations in a serverless environment. Azure Service Bus is a fully managed message broker service that enables reliable and secure communication between applications. Azure Functions integrates with Azure Service Bus via triggers and bindings, allowing you to build functions that react to and send queue or topic messages.

upvoted 1 times

🗨️ 👤 **FeriAZ** 10 months, 3 weeks ago

**Selected Answer: A**

By leveraging an Azure Service Bus queue to decouple the initial HTTP request from the blob data processing, this architectural pattern effectively addresses the timeout issue. It ensures that the app can process blob data without being constrained by the execution timeout of the HTTP-triggered function, aligning well with best practices for building scalable and resilient cloud applications.

upvoted 1 times

🗨️ 👤 **Jak007** 11 months ago

**Selected Answer: A**

A) Yes. This is true and is a suggested approach from the documentation here: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

upvoted 1 times

🗨️ 👤 **garbas** 11 months, 1 week ago

**Selected Answer: B**

If you return immediate success, the app can't ensure the blob data is actually processed

<https://learn.microsoft.com/en-us/azure/azure-functions/performance-reliability#make-sure-background-tasks-complete>

upvoted 2 times

🗨️ 👤 **bgbgvfvf** 1 year ago

A yes is correct

upvoted 1 times

🗨️ 👤 **AlbertoBT** 1 year, 1 month ago

**Selected Answer: B**

YES but only to ensure the app does not time out but not to processes the blob data

An HTTP triggered function has a 320 seconds of max timeout and is because the HTTP request binding. So the solution only avoids having a timeout in the request but not in the processing

The max timeout for any function is 4 minutes (With another kind of triggers) so the only way to be sure that the blob is processed is with a durable function

The right answer is NO

upvoted 1 times

🗨️ 👤 **mlop3s** 1 year, 1 month ago

The problem here is "The app must process the blob data.". Using a queue trigger will not guarantee that the blob is processed. Asynchronous pattern will let me ask for the status and if it fails, I can resend the request. I would go for no.

upvoted 1 times

🗨️ 👤 **zreaf29** 1 year, 1 month ago

I think answer is B.No cause the Http trigger can take more time than four minute and it cause time out. Best way is using durable function which can check stateful

upvoted 1 times

🗨️ 👤 **Vladimir\_Gajinov** 1 year, 3 months ago

**Selected Answer: A**

The answer is YES.

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions

async pattern or defer the actual work and return an immediate response.

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

upvoted 2 times

🗨️ 👤 **nardk** 1 year, 3 months ago

**Selected Answer: A**

Answer is A

upvoted 1 times

🗨️ 👤 **Kanasan** 1 year, 4 months ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response.

upvoted 1 times

🗨️ 👤 **longnguyendh** 1 year, 5 months ago

A is correct.

upvoted 1 times



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 develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob. The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Configure the app to use an App Service hosting plan and enable the Always On setting.

Does the solution meet the goal?

- A. Yes
- B. No

**Suggested Answer: B**

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.


Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

Community vote distribution



**00avatar** Highly Voted 3 years, 11 months ago

Answer "No" is correct. Always On enables waking up on HTTPTrigger, but does not prevent the exceeding the max time out time of 230 seconds.  
<https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan#always-on>  
 upvoted 65 times

**Knightie** 2 years, 4 months ago

No is the answer, but the reason is the timeout being raised by HTTP layer from the Azure Load Balancer, not the App layer that at least it gives 5 minutes for the cheapest type, Consumption, so however you enhance the app layer, the http layer Azure Load Balance will not wait more than 230 second and will reply it as timeout. Use the durable function pattern to poll the status for completion will be the easiest solution, else avoid the http layer like service bus will work too.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout~:text=1%20Regardless%20of,immediate%20response>.  
 upvoted 6 times

**surprise0011** 1 year, 8 months ago  
 received 2023-04-17 went given answer, score 926  
 upvoted 2 times

**abdou1987** 3 years, 7 months ago  
 Answer is YES  
 the default time for App service plan is 30 to unlimit  
<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>  
 upvoted 3 times

**abdou1987** 3 years, 7 months ago  
 Sorry i cant delete my previois comment.  
 I confirm the answer is NO. Whatever the app service plan the maximum timeout for HTTP triggers is 230 seconds.  
 upvoted 23 times

**Amrit862** 3 years, 11 months ago

Answer is still 'No' but default time for dedicated host is 30 mins, it can be technically configured to infinite...but have to do that manually and that is not mentioned in solution.

ref: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-host-json#functiontimeout>

upvoted 1 times

 **Cornholioz** Highly Voted 4 years, 1 month ago

Answer "B.No" seems right because this is not about App Services Best Practices.


upvoted 9 times

 **VK1989** Most Recent 5 months, 3 weeks ago

Selected Answer: B

Answer is "No" just verified with GEMINI

upvoted 2 times

 **HariB1992** 8 months, 1 week ago

CoPilot Says Option as Yes.

Yes, the solution meets the goal. By configuring the app to use an App Service hosting plan and enabling the Always On setting, you can prevent the app from timing out and ensure that it processes the blob data successfully.


upvoted 1 times

 **[Removed]** 1 year, 1 month ago

Selected Answer: B

No i correct

upvoted 1 times


 **Kanasan** 1 year, 4 months ago

Selected Answer: B

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer.


upvoted 1 times

 **RenatoJr** 1 year, 8 months ago

Selected Answer: B


Answer "No" is correct

upvoted 1 times

 **Anitay** 2 years, 3 months ago

Use Durable function for this requirement

upvoted 1 times

 **Perplex** 2 years, 4 months ago

Selected Answer: B

Always On does not prevent a timeout so is not a solution to this problem. Answer is B (No).


upvoted 1 times

 **Rakeshpro** 2 years, 5 months ago

Selected Answer: A

Correct!

upvoted 1 times

 **Eltooth** 2 years, 6 months ago

Selected Answer: B

B is correct answer (No).

Max timeout for HTTP triggers is 230 seconds < 240 seconds (4 minutes).


upvoted 1 times

 **TakumaK** 3 years, 6 months ago

"Always On" feature of Azure App Service is to keep the host process running to allow more responsive to requests after significant idle periods. This is UNRELATED to the timeout.

So the answer is obvious!!

upvoted 5 times

 **mlantonis** 3 years, 7 months ago

Correct Answer: No

Always On enables waking up on HTTP trigger, but does not prevent the exceeding the max time out time of 230 seconds.

If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. On an App Service plan, the functions runtime goes idle after a few minutes of inactivity, so only HTTP triggers will "wake up" your functions. The Always on setting is available only on an App Service plan. On a Consumption plan, the platform activates function apps automatically.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/dedicated-plan#always-on>  
upvoted 8 times

  **abdou1987** 3 years, 7 months ago

Answer is YES

"Best for long-running scenarios where Durable Functions can't be used. Consider an App Service plan in the following situations:" link:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#overview-of-plans>

another link <https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout> that mentionned the duration is 30 min and upper  
upvoted 1 times

  **abdou1987** 3 years, 7 months ago

Sorry i cant delete my previois comment.

I confirm the answer is NO. Whatever the app service plan the maximum timeout for HTTP triggers is 230 seconds.

upvoted 3 times

  **TakumaK** 3 years, 6 months ago

you are messing up this comment section by correcting your own wrong answer. :)

upvoted 1 times

  **UnknowMan** 3 years, 7 months ago

Answer "No" is correct

upvoted 2 times

  **glam** 3 years, 7 months ago

B. No,,

upvoted 3 times

  **bugimachi** 4 years ago

The page, however, does not talk about Service Bus queues (any more).

upvoted 1 times

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure

Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Move photo processing to an Azure Function triggered from the blob upload.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: A

Azure Storage events allow applications to react to events. Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow.


Events are pushed using Azure Event Grid to subscribers such as Azure Functions, Azure Logic Apps, or even to your own http listener.

Note: Only storage accounts of kind StorageV2 (general purpose v2) and BlobStorage support event integration. Storage (general purpose v1) does not support integration with Event Grid.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

Community vote distribution



**AnkanG** Highly Voted 4 years, 11 months ago

as is correct

upvoted 57 times

**fadikh** Highly Voted 4 years, 4 months ago

The answer is correct, but it should have also mentioned that the function app must not be on a consumption plan. because in that case, it might take up to 10 minutes to process the event.

upvoted 31 times

**BrettusMaximus** 4 years, 1 month ago

So the answer is NO. Cant wait 10 minutes

upvoted 8 times

**altafpatel1984** 3 years, 7 months ago

That limitation is only for Consumption plan. Since no mention of it, we can assume user can opt for premium, app service plan.

upvoted 5 times

**HelloCloudWorld** 1 year ago

Or we can assume user doesn't use premium plan... your argument isn't valid

upvoted 1 times

**bbou** 4 years, 2 months ago

No 10 seconds

upvoted 1 times

**bbou** 4 years, 2 months ago

I was wrong - yes up to 10 minutes

upvoted 6 times

**dookebroyle** 3 years ago

The time limit is for the process to START only. It can take longer to process.

upvoted 6 times

🗨️ 👤 **Ody** Most Recent 5 months, 3 weeks ago

**Selected Answer: A**

It appears this answer has changed over time and now there are two ways to configure the trigger.

Newer --> Event-Based Trigger: This implementation uses Azure Event Grid to trigger the function. It provides low latency, typically triggering the function within seconds of the blob upload.

Older --> Polling-Based Trigger: This implementation periodically polls the storage account for changes. The latency can be higher, potentially up to several minutes, depending on the polling interval.

I only question if anyone writing test questions at Microsoft knows about it.

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-event-grid-blob-trigger?pivot=programming-language-csharp>  
upvoted 2 times

🗨️ 👤 **Tej\_007** 11 months ago

**Selected Answer: B**

blob storage event doesn't guarantee an SLA.

you cannot control the event arrival in less than a minute.

upvoted 2 times

🗨️ 👤 **Munwalinwali** 12 months ago

**Selected Answer: A**

Answer is yes. Posted the question in ChatGPT

upvoted 1 times

🗨️ 👤 **Isoldhe** 1 year, 4 months ago

**Selected Answer: A**

Copilot AI:

Yes, the solution meets the goal. Moving photo processing to an Azure Function triggered from the blob upload can help you create a scalable, serverless, and event-driven solution for your SaaS offering. Azure Functions can run code in response to various triggers, such as HTTP requests, timers, queues, and blobs. In this case, you can use a Blob storage trigger to execute your photo processing logic whenever a new photo is uploaded to a specific container in your storage account. This way, you can ensure that the process to produce a mobile-friendly version of the image starts in less than one minute, as the function is invoked almost immediately after the blob is created or updated.

upvoted 2 times

🗨️ 👤 **ragha81** 1 year, 4 months ago

Answer is B as per below comment

There is no service level agreement around the time it takes for a message to arrive. It's not uncommon for messages to arrive anywhere from 30 minutes to two hours.

upvoted 2 times

🗨️ 👤 **AhmedAbdelAziz** 1 year, 5 months ago

Correct No : As there is no guarantee on the blob storage to trigger in less than min , so we should use the event grid and the azure function should subscribe to the event grid not the blob direct

upvoted 2 times

🗨️ 👤 **Jak007** 1 year, 5 months ago

**Selected Answer: B**

B) No.

As per the documentation linked by @Rabi3. Under "Recommended Practices for Consuming Blob Events": "There is no service level agreement around the time it takes for a message to arrive. It's not uncommon for messages to arrive anywhere from 30 minutes to two hours".

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview#practices-for-consuming-events>

upvoted 2 times

🗨️ 👤 **Rabi3** 1 year, 5 months ago

**Selected Answer: B**

Answer is B. the key point here is that the processing should start in less one minute. Which is not a guaranteed for Blob storage events.

"There is no service level agreement around the time it takes for a message to arrive. It's not uncommon for messages to arrive anywhere from 30 minutes to two hours."

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview#practices-for-consuming-events>

upvoted 2 times

🗲️ 👤 **bgbgvf** 1 year, 6 months ago

I think the answer is B

upvoted 1 times

🗲️ 👤 **ami2023** 1 year, 7 months ago

**Selected Answer: A**

this is related to topic 2 question 8, A. Yes should the correct answer.

upvoted 1 times

🗲️ 👤 **dddddd111** 1 year, 8 months ago

**Selected Answer: A**

I think the answer is A.

upvoted 1 times

🗲️ 👤 **Dixavado** 1 year, 9 months ago

**Selected Answer: A**

Go with A

upvoted 1 times

🗲️ 👤 **Kanasan** 1 year, 10 months ago

> The process to produce a mobile-friendly version of the image must start in less than one minute.

This cannot be fulfilled with the proposed solution.

So, vote for B

upvoted 1 times

🗲️ 👤 **LxstlnNxght** 1 year, 10 months ago

**Selected Answer: A**

Answer is A

upvoted 1 times

🗲️ 👤 **EliteAllen** 1 year, 11 months ago

**Selected Answer: A**

A. Yes

The solution does meet the goal. Azure Functions can be triggered by blob storage events, such as the uploading of a new photo. When a new photo is uploaded to blob storage, the Azure Function would be triggered to process the photo and create a mobile-friendly version. The trigger is almost instantaneous, so the process would indeed start in less than a minute after a photo is uploaded. This meets the requirement mentioned in the question.

upvoted 1 times

You are developing an application that uses Azure Blob storage.

The application must read the transaction logs of all the changes that occur to the blobs and the blob metadata in the storage account for auditing purposes. The changes must be in the order in which they occurred, include only create, update, delete, and copy operations and be retained for compliance reasons.

You need to process the transaction logs asynchronously.

What should you do?

- A. Process all Azure Blob storage events by using Azure Event Grid with a subscriber Azure Function app.
- B. Enable the change feed on the storage account and process all changes for available events.
- C. Process all Azure Storage Analytics logs for successful blob events.
- D. Use the Azure Monitor HTTP Data Collector API and scan the request body for successful blob events.

**Suggested Answer: B**

Change feed support in Azure Blob Storage

The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

Community vote distribution

B (100%)

 **Ummara** Highly Voted 4 years, 11 months ago


B: the change feed provides transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, a durable, immutable, read-only log of these changes. You can process these logs asynchronously, incrementally or in-full.

upvoted 80 times

 **azurelearner666** 4 years ago

Right!

upvoted 3 times

 **mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: B

The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account. The change feed provides ordered, guaranteed, durable, immutable, read-only log of these changes. Client applications can read these logs at any time, either in streaming or in batch mode. The change feed enables you to build efficient and scalable solutions that process change events that occur in your Blob Storage account at a low cost.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed>

upvoted 20 times

 **CESBCN** Most Recent 8 months, 2 weeks ago

**Selected Answer: B**

b

upvoted 1 times

 **darcimf** 11 months ago

I think B is correct now, but I see copy operation and guessed that's not a change.

upvoted 1 times

🗨️ 👤 **Sanjeev5131** 1 year, 1 month ago

Ans --> A

Enabling the change feed on the storage account) might not be the ideal choice for asynchronously processing transaction logs of blob changes for auditing purposes, compared to using Azure Event Grid:

Change Feed vs. Azure Event Grid:

Enabling the change feed on the storage account provides a log of blob changes, but it's designed more for data replication and integration scenarios rather than real-time event processing.

Azure Event Grid, on the other hand, is specifically tailored for event-driven architectures and can deliver near real-time notifications about blob operations, which is crucial for auditing purposes where timeliness and reliability of events matter.

upvoted 2 times

🗨️ 👤 **Jak007** 1 year, 5 months ago

**Selected Answer: B**

B) As stated by Exam topics.

They also provide the link to the Learn resource. Highly recommend reading it. One of the best Learn articles I've come across!

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=template>

upvoted 2 times

🗨️ 👤 **deathRac3** 2 years, 1 month ago

**Selected Answer: B**

as is correct

upvoted 2 times

🗨️ 👤 **Saluk\_DE** 2 years, 3 months ago

Question was in Exam 2023-03-30

upvoted 4 times

🗨️ 👤 **NombreFalso** 2 years, 4 months ago

**Selected Answer: B**

Exactly what Change Feed offers

upvoted 1 times

🗨️ 👤 **alexein74** 2 years, 5 months ago

**Selected Answer: B**

B. Enable the change feed on the storage account and process all changes for available events.

Azure Blob storage change feed provides a log of all the create, update, delete, and copy operations that occur on blobs and blob metadata in the storage account. It allows you to track the changes to the blobs in the order in which they occurred, which is what you need for auditing purposes. By enabling the change feed, your application can asynchronously process the changes and retain them for compliance reasons. You can also use Azure Event Grid to route the change feed events to a subscriber Azure Function app for further processing, this way you can keep the logs for the compliance reasons.

upvoted 1 times

🗨️ 👤 **carlosghosn** 2 years, 7 months ago

Got this in the exam today ! Nov 25, 2022

upvoted 2 times

🗨️ 👤 **AZAdam22** 2 years, 11 months ago

**Selected Answer: B**

B - Because change feed contains an ordered list of operations that have been made in a storage account.

upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

**Selected Answer: B**

B is correct answer.

upvoted 1 times

🗨️ 👤 **hamzabts** 3 years, 1 month ago

B : The purpose of the change feed is to provide transaction logs of all the changes that occur to the blobs and the blob metadata in your storage account.

upvoted 1 times

🗨️ 👤 **Freidrich** 3 years, 4 months ago



Selected Answer: B

Correct.

upvoted 1 times

🗨️ 👤 **HimanshuNankani** 3 years, 8 months ago

What about the delete operations that are not logged by the change feed?

upvoted 6 times

🗨️ 👤 **santoshsidnal** 3 years ago

That's a valid question!! I think we should use event grid approach i.e A

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 9 months ago

BlobDeleted is supported eventType and even if it is not supported still the requirement that it should be ordered cannot be satisfied with the blob-storage-events through event-grid. Change feed supports ordering

upvoted 1 times

🗨️ 👤 **ozbonny** 3 years, 11 months ago

enabled change feed to get information and for auditing or compliance purpose

upvoted 1 times

## DRAG DROP -

You plan to create a Docker image that runs an ASP.NET Core application named ContosoApp. You have a setup script named setupScript.ps1 and a series of application files including ContosoApp.dll.

You need to create a Dockerfile document that meets the following requirements:

- ⇒ Call setupScripts.ps1 when the container is built.
- ⇒ Run ContosoApp.dll when the container starts.

The Dockerfile document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.

Which five commands should you use to develop the solution? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

### Commands

FROM microsoft/aspnetcore:latest

WORKDIR /apps/ContosoApp

CMD ["dotnet", "ContosoApp.dll"]

COPY ./ .

RUN powershell ./setupScript.ps1

### Answer Area

Suggested Answer:

### Commands

FROM microsoft/aspnetcore:latest

WORKDIR /apps/ContosoApp

CMD ["dotnet", "ContosoApp.dll"]

COPY ./ .

RUN powershell ./setupScript.ps1

### Answer Area

CMD ["dotnet", "ContosoApp.dll"]

FROM microsoft/aspnetcore:latest

WORKDIR /apps/ContosoApp

COPY ./ .

RUN powershell ./setupScript.ps1

Box 1: CMD [..]

Cmd starts a new instance of the command interpreter, Cmd.exe.

Syntax: CMD <string>

Specifies the command you want to carry out.

Box 2: FROM microsoft/aspnetcore-build:latest

Box 3: WORKDIR /apps/ContosoApp -

Bxo 4: COPY ./ .

Box 5: RUN powershell ./setupScript.ps1

 **agueda** Highly Voted 4 years, 3 months ago

It should be:

- FROM
- WORKDIR
- COPY
- RUN
- CMD

Same question on:

<https://www.examttopics.com/discussions/microsoft/view/13131-exam-az-300-topic-3-question-4-discussion/>

And:

<https://www.examttopics.com/discussions/microsoft/view/11045-exam-az-203-topic-1-question-7-discussion/>

upvoted 303 times

🗨️ 👤 **PereCastor37** 6 months, 2 weeks ago

Agree with you :

Dockerfile example by chatgpt 4o without give him the sentence :

FROM

WORKDIR

COPY

RUN

CMD

upvoted 1 times

🗨️ 👤 **Den1354** 3 years, 5 months ago

- FROM

- COPY

- WORKDIR

- RUN

- CMD

Otherwise we are going to set for work directory path which doesn't exist yet

upvoted 5 times

🗨️ 👤 **zolani** 1 year, 2 months ago

True

FROM microsoft/aspnetcore:latest

This command specifies the base image for your container.

WORKDIR /apps/ContosoApp

This command sets the working directory for the container.

COPY ./ .

This command copies all files and folders from the current directory (where the Dockerfile resides) to the working directory (/apps/ContosoApp) inside the container.

RUN powershell ./setupScript.ps1

This command runs the setupScript.ps1 script using PowerShell within the container during the build process.

CMD ["dotnet", "ContosoApp.dll"]

This command defines the default command that will be executed when the container starts.

upvoted 4 times

🗨️ 👤 **Dinima** 4 years, 3 months ago

You are correct. This has been discussed in Udemy course as well as follows,

The first statement in the Dockerfile must be the FROM statement to specify the image to use as the base image.

Then specify the Image working directory

Then copy all of the application contents using the COPY command

And then use the CMD command to run the PowerShell command and the ENTRYPOINT statement to run the dotnet application.

upvoted 29 times

🗨️ 👤 **Basu525** 4 years, 3 months ago

as per Udemy, the last steps would be CMD powershell ./script.ps1 and then ENTRYPOINT (dotnet, xx.dll) which I believe is the correct answer.

But unfortunately the options are not there in Examtopic

upvoted 3 times

🗨️ 👤 **ranjitklive** 3 years, 11 months ago

ENTRYPOINT instruction works very similarly to CMD in that it is used to specify the command executed when the container is started.

upvoted 3 times

🗨️ 👤 **solidrock** 3 years, 8 months ago

which udemy course you guys are talking about?

upvoted 2 times

🗨️ 👤 **balis** 3 years, 3 months ago

This is correct answer

It should be:

- FROM
- WORKDIR
- COPY
- RUN
- CMD

because WORKDIR will create directory if it doesn't exist <https://docs.docker.com/engine/reference/builder/#workdir>

upvoted 14 times

🗨️ 👤 **TakumaK** Highly Voted 👍 4 years ago

Just wondering who put the answers for the questions in this site? most of them are not correct.

upvoted 25 times

🗨️ 👤 **SlavMar** 4 years ago

VCE exams have same issues

Probably this are some braindumps of people who were not well prepared to take test or they too test just to scrap exam questions but they have no matter knowledge

upvoted 4 times

🗨️ 👤 **heptadecane** Most Recent 🔍 7 months, 3 weeks ago

FROM mcr.microsoft.com/dotnet/aspnet:latest

WORKDIR /apps/ContosoApp

COPY ./ .

RUN powershell ./setupScript.ps1

CMD ["dotnet", "ContosoApp.dll"]

upvoted 3 times

🗨️ 👤 **4bd3116** 10 months, 1 week ago

FROM microsoft/aspnetcore:latest

COPY . /apps/ContosoApp

WORKDIR /apps/ContosoApp

RUN powershell ./setupScript.ps1

CMD ["dotnet", "ContosoApp.dll"]

upvoted 1 times

🗨️ 👤 **ciamp** 10 months, 2 weeks ago

FROM ubuntu:latest

WORKDIR /app

COPY . /app

RUN apt-get update && apt-get install -y <dependencies>

CMD ["executable", "param1", "param2"]

upvoted 1 times

🗨️ 👤 **onlyforheros** 1 year, 3 months ago

Got it in the exam 13.03.2024. Score: 910. Went with

- FROM
- WORKDIR
- COPY
- RUN
- CMD

upvoted 3 times

🗨️ 👤 **AhmedAbdelAziz** 1 year, 5 months ago

Answer is

- FROM
- RUN
- WORKDIR
- COPY
- CMD

check this link

<https://learn.microsoft.com/en-us/training/modules/intro-to-docker-containers/3-how-docker-images-work>

upvoted 3 times

  **manopeyda** 1 year, 5 months ago

Here are the five commands you can use in the Dockerfile:

FROM: Specify the base image.

Dockerfile

Copy code

FROM mcr.microsoft.com/dotnet/core/aspnet:3.1

COPY: Copy the application files to the container.

Dockerfile

Copy code

COPY . /app

WORKDIR: Set the working directory to the application folder.

Dockerfile

Copy code

WORKDIR /app

RUN: Execute the setup script during the image build process.

Dockerfile

Copy code

RUN ./setupScript.ps1

CMD: Specify the command to run when the container starts.


Dockerfile

Copy code

CMD ["dotnet", "ContosoApp.dll"]



Make sure to adjust the paths and filenames according to your actual file structure.

upvoted 2 times

  **arunkum** 1 year, 6 months ago

Got it in the exam 14/12/23. All questions are from ExamTopics. Case study - VanArsdel, Ltd (11 questions)

upvoted 1 times

  **p2006** 1 year, 9 months ago

Got on 9/25/2023

From


Workdir

Copy

Run powershell


Cmd ContosoApp.dll

upvoted 3 times

  **RuffBoii** 1 year, 9 months ago

Had this on my exam today.

upvoted 1 times

  **mihailos** 1 year, 10 months ago

Got it in exam 28/08/23. Scored 912. Went with the following:

- FROM



- WORKDIR

- COPY

- RUN

- CMD

upvoted 1 times

  **TechyNetty** 1 year, 10 months ago

Im yet to appear for exam soon. Since you attempted recently, were most of the questions from this site? Please respond, it will be helpful. thanks  
upvoted 1 times

🗨️ 👤 **longnguyendh** 1 year, 11 months ago

- FROM ...
- WORKDIR ...
- COPY ...
- RUN ...
- CMD ...

upvoted 1 times

🗨️ 👤 **deathRac3** 2 years ago

Question was there for me on 29th May 2023

upvoted 1 times

🗨️ 👤 **70PineApple** 2 years, 4 months ago

Got this in exam today..20/02/23

score: 817

upvoted 1 times

🗨️ 👤 **Priya0703** 2 years, 4 months ago

Got this question today on 20-02-2023 exam.

upvoted 1 times

🗨️ 👤 **HafizSalmanMalik** 2 years, 5 months ago

It should be:

- FROM
- WORKDIR
- COPY
- RUN
- CMD

upvoted 3 times

You are developing an Azure Function App that processes images that are uploaded to an Azure Blob container. Images must be processed as quickly as possible after they are uploaded, and the solution must minimize latency. You create code to process images when the Function App is triggered. You need to configure the Function App. What should you do?

- A. Use an App Service plan. Configure the Function App to use an Azure Blob Storage input trigger.
- B. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage trigger.
- C. Use a Consumption plan. Configure the Function App to use a Timer trigger.
- D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.
- E. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage input trigger.


**Suggested Answer: B**

The Blob storage trigger starts a function when a new or updated blob is detected. The blob contents are provided as input to the function. The Consumption plan limits a function app on one virtual machine (VM) to 1.5 GB of memory.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger>

Community vote distribution



**Kitkit** Highly Voted 4 years, 4 months ago

The answer is D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger. Consumption plan can cause a 10-min delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp>

upvoted 205 times

**SlavMar** 4 years ago

Why not B then.

Using input for function apps seems to make more sense

upvoted 4 times

**zzt** 3 years, 11 months ago

Seemingly there is no such thing as "input trigger" it is "input binding"

See: <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-input?tabs=csharp>

upvoted 35 times

**SlavMar** 4 years ago

I ment A

upvoted 1 times

**nosby** 1 year, 2 months ago

Agree. <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#cold-start-behavior>

upvoted 1 times

**samraw83** 3 years ago

App service plan is best for the long running scenarios where you cannot use durable functions.

upvoted 2 times

**OscarL** 1 year, 12 months ago

totally agree

upvoted 1 times

**danielcr** Highly Voted 4 years, 4 months ago

A or D, never B.

With Consumption plan, you could have cold start, and the question say "must be processed as quickly as possible" so you need an App Service Plan.


Between A and D ... All triggers are Input, so i don't understand the difference.

upvoted 24 times

  **ranjitklive** 3 years, 11 months ago

Agree with your point, all triggers are input...!!

upvoted 3 times

  **ewertonews** 3 years, 9 months ago

But that nomenclature is not used by MS.

A function can have a trigger and an input binding (HTTP trigger with a Storage Queue as input binding for example).

upvoted 3 times

  **ewertonews** 3 years, 9 months ago

It cannot be A. There is no such thing as an input trigger. It's either a (Blob Storage) trigger or an input binding. They are different things.

upvoted 13 times

  **surprise0011** 2 years, 2 months ago

Exactly my friend. I was trying to find info about it and as I was thinking, there is no such thing



upvoted 1 times

  **laminall** Most Recent 4 months ago

**Selected Answer: D**

Input binding not input trigger, it doesn't exist

upvoted 1 times

  **bouda19** 5 months, 2 weeks ago

**Selected Answer: D**

Response of copilot:

To minimize both latency and cost for uploading media to Azure Blob Storage, the Premium Plan would be the best choice. Here's why:


Latency: The Premium Plan eliminates cold starts by keeping instances warm, which means your function will respond faster when triggered.

Cost: While the Premium Plan has a higher base cost compared to the Consumption Plan, it offers more predictable pricing and can be more cost-effective if your function runs continuously or frequently.

The Consumption Plan is cost-effective for sporadic workloads but may suffer from cold starts, leading to higher latency. The Service Plan provides more control and can be kept "Always On," but it might be more expensive for infrequent use.

So the answer is D

upvoted 1 times

  **DD7826** 5 months, 3 weeks ago

**Selected Answer: B**

To process images as quickly as possible after they are uploaded to an Azure Blob container and minimize latency, you should:

B. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage trigger.

This configuration ensures that your Function App is triggered immediately when a new image is uploaded to the Blob container, leveraging the serverless nature of the Consumption plan to scale automatically and handle the processing efficiently

upvoted 1 times

  **ns4098** 9 months ago

B. Use a Consumption plan. Configure the Function App to use an Azure Blob Storage trigger.

Here's why:

Consumption Plan: This plan is optimized for event-driven tasks like image processing, and it automatically scales based on the number of incoming requests, ensuring that images are processed as quickly as possible. The Function App will scale to handle multiple images without pre-allocating resources.

Azure Blob Storage Trigger: This trigger is designed to respond to new blobs (images) in a storage container. It will fire the function as soon as a new blob is uploaded, which minimizes latency and ensures images are processed immediately.

Other options:



App Service Plan: Not the best choice for minimizing latency and cost. It requires pre-allocated resources, which might lead to higher costs and unnecessary resource allocation, especially if your function is event-driven and sporadic.

Timer Trigger: This would introduce delays, as the function would be triggered on a schedule, not immediately after a new image is uploaded.

upvoted 1 times

🗨️ 👤 **Faolba** 9 months, 1 week ago

**Selected Answer: B**

B- Using a Consumption plan with an Azure Blob Storage trigger is the most effective way to ensure that images are processed as quickly as possible after they are uploaded. The Blob Storage trigger automatically invokes the Function App in response to blob uploads, minimizing latency and handling scaling automatically based on demand.

While option D is technically correct, using an App Service plan may not be the most efficient choice for this scenario if your goal is to minimize latency and optimize costs.

A Consumption plan is typically more suitable for event-driven scenarios like this, as it automatically scales and charges based on usage. An App Service plan might provide more predictable performance but could be more costly if the function is not in constant use. Therefore, option B remains the more optimal choice for processing images immediately upon upload.

upvoted 1 times

🗨️ 👤 **macfuk** 1 year ago

Whenever you see Function App, then answer should contain the word "consumption plan"

upvoted 1 times

🗨️ 👤 **Stel0Papad4** 1 year, 1 month ago

**Selected Answer: D**

After a short discussion with Copilot AI , it gave me this answer:

(Answer)

If performance and no latency are top priorities, go with Option D (App Service plan).

If cost efficiency and event-driven scalability matter more, choose Option B (Consumption plan).

upvoted 3 times

🗨️ 👤 **Isoldhe** 1 year, 4 months ago

**Selected Answer: B**

Copilot AI also says the answer is B

upvoted 1 times

🗨️ 👤 **TheFivePips** 10 months, 1 week ago

AI has a hard time with details some of the time. If you remind it that consumption plans may take up to 10 minutes to warm up then it will correct itself

upvoted 1 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

Got this answer from ChatGPT:

I understand the concern for minimizing latency and processing images as quickly as possible. However, the Azure Functions Consumption plan is designed to scale automatically based on demand, providing a cost-effective and efficient solution for scenarios where latency is a critical factor.

When you configure your Azure Function App with an Azure Blob Storage trigger in a Consumption plan, it enables the function to respond quickly to new blob uploads. The Consumption plan is optimized for event-driven workloads, making it suitable for scenarios where you need fast response times.

Therefore, option B (Use a Consumption plan. Configure the Function App to use an Azure Blob Storage trigger) remains a suitable and efficient choice for your requirements.

upvoted 3 times

🗨️ 👤 **RuffBoii** 1 year, 9 months ago

Had this on my exam today.

upvoted 3 times

🗨️ 👤 **CSLK** 1 year, 3 months ago

What the answer you gave

upvoted 2 times

🗨️ 👤 **Guurr40\_\_** 1 year, 9 months ago

What is the difference between "storage trigger" and "storage input trigger"? In the end, isn't the second one just a reinforcement to indicate that the moment a file is placed, precisely as input, in the hypothetical folder, then it is picked up and processed by the trigger. Thank you in advance for your answers

upvoted 1 times

🗨️ 👤 **sachinrikhe** 2 years, 3 months ago

Answer is B.

When using a Consumption plan, Azure automatically scales out instances of the function as needed to handle incoming requests, and you only pay for the actual number of executions of the function. This ensures that the function is always available to process images without any delays caused by having to scale up or down manually.

upvoted 5 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

You can't use Consumption Plan. Consumption Plan can take up to 10 mins to warm up which doesn't fit the requirement.

upvoted 2 times

🗨️ 👤 **nedlo** 2 years, 3 months ago

**Selected Answer: D**

D is correct "For this scenario where images must be processed as quickly as possible and latency must be minimized, it is recommended to use the App Service plan instead of the Consumption plan.

The App Service plan provides dedicated resources to the function app, which can help improve performance and minimize the risk of cold start delays. Additionally, the App Service plan allows you to scale out the function app to multiple instances, which can help handle high traffic loads and further improve performance.

On the other hand, the Consumption plan is a serverless hosting option that scales automatically based on demand, but may result in cold start delays and limited resources during high traffic periods.

Therefore, it is recommended to use the App Service plan for this scenario."

upvoted 3 times

🗨️ 👤 **vguimars** 2 years, 4 months ago

**Selected Answer: D**

Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.

Consumption plan can cause a 10-min delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp>

upvoted 1 times

🗨️ 👤 **OPT\_001122** 2 years, 6 months ago

**Selected Answer: D**

D. Use an App Service plan. Configure the Function App to use an Azure Blob Storage trigger.

upvoted 1 times

## HOTSPOT -

You are configuring a new development environment for a Java application.

The environment requires a Virtual Machine Scale Set (VMSS), several storage accounts, and networking components.

The VMSS must not be created until the storage accounts have been successfully created and an associated load balancer and virtual network is configured.

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
{
  ...
  "resources": [
    {
      "apiVersion": "2016-01-01",
      "type": "Microsoft.Storage/storageAccounts",
      "name": "[concat(
        ( ), 'storage', uniqueString(resourceGroup().id))]",
      "location": "[resourceGroup().location]",
      "sku": {
        "name": "Standard_LRS"
      },
      "kind": "Storage",
      "properties": {},
      "tags": {
        "copy": " ",
        "copyIndex": " ",
        "priority": " ",
        "dependsOn": " "
      },
      "name": "storagesetup",
      "count": 3
    },
    {
      "apiVersion": "2015-06-15",
      "type": "Microsoft.Compute/virtualMachines",
      "name": "[concat('VM', uniqueString(resourceGroup().id))]",
      "tags": {
        "copy": " ",
        "copyIndex": " ",
        "priority": " ",
        "dependsOn": " "
      },
      "[variables('loadBalancerName')]",
      "[variables('virtualNetworkName')]",
      "storagesetup",
    },
    ...
  ],
  "outputs": {}
}
```

Suggested Answer:

### Answer Area

```
{
  . . .
  "resources": [
    {
      "apiVersion": "2016-01-01",
      "type": "Microsoft.Storage/storageAccounts",
      "name": "[concat(
        

|           |
|-----------|
|           |
| copy      |
| copyIndex |
| priority  |
| dependsOn |


        ), 'storage', uniqueString(resourceGroup().id))]",

      "location": "[resourceGroup().location]",
      . . .
      "sku": {
        "name": "Standard_LRS"
      },
      "kind": "Storage",
      "properties": {},
      "
        

|           |
|-----------|
|           |
| copy      |
| copyIndex |
| priority  |
| dependsOn |


        ": {
          "name": "storagesetup",
          "count": 3
        }
    },
    {
      "apiVersion": "2015-06-15",
      "type": "Microsoft.Compute/virtualMachines",
      "name": "[concat('VM', uniqueString(resourceGroup().id))]",
      "
        

|           |
|-----------|
|           |
| copy      |
| copyIndex |
| priority  |
| dependsOn |


        ": [
          "[variables('loadBalancerName')]",
          "[variables('virtualNetworkName')]",
          "storagesetup",
        ],
      . . .
    }
  ],
  "outputs": {}
}
```

Box 1: copyIndex -

Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.

Box 2: copy -

By adding the copy element to the resources section of your template, you can dynamically set the number of resources to deploy.

Box 3: dependsOn -

Example:

```
"type": "Microsoft.Compute/virtualMachineScaleSets",
"apiVersion": "2020-06-01",
"name": "[variables('namingInfix')]",
"location": "[parameters('location')]",
"sku": {
  "name": "[parameters('vmSku')]",
  "tier": "Standard",
  "capacity": "[parameters('instanceCount')]"
},
"dependsOn": [
  "[resourceId('Microsoft.Network/loadBalancers', variables('loadBalancerName'))]",
```

```
"[resourceId('Microsoft.Network/virtualNetworks', variables('virtualNetworkName'))]"
```

```
],
```



Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources> <https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows>

  **cloudlabadm** Highly Voted 4 years, 3 months ago

Answer: copyIndex, copy, Dependson. Check this link <https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/define-resource-dependency>.

upvoted 80 times

  **Esward** 2 years, 5 months ago

Agreed

upvoted 3 times

  **130nk3r5** 1 year, 6 months ago

Got this today.

Went with answer here.

Score 927

upvoted 5 times

  **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: copyIndex

Notice that the name of each resource includes the copyIndex() function, which returns the current iteration in the loop. copyIndex() is zero-based.

Box 2: copy

By adding copy loop to the resources section of your template, you can dynamically set the number of resources to deploy. You also avoid having to repeat template syntax.

Box 3: dependsOn

Within your Azure Resource Manager template (ARM template), the dependsOn element enables you to define one resource as a dependent on one or more resources.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/copy-resources>

<https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/quick-create-template-windows>

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/define-resource-dependency>

upvoted 57 times

  **ben\_thedevguy** Most Recent 9 months, 4 weeks ago

Arm template is no longer in the AZ-204 exam according to this answer:

<https://trainingsupport.microsoft.com/en-us/mcp/forum/all/is-azure-resource-manager-arm-in-the-az-204-exam/bcdfdec6-2d40-470a-ba24-e3e6de9fdfaa>



Although I have not yet taken the exam since the answer is from a Microsoft Moderator which I think is reliable.

upvoted 1 times

  **bgbgvfvf** 1 year, 6 months ago

I think answer is Correct!

upvoted 1 times

  **MrAzz** 1 year, 8 months ago

For the second part: "Copy"

What is the point of adding the same property with the same name 3 times?

upvoted 1 times

  **deathRac3** 2 years ago

Question was there for me on 29th May 2023

upvoted 3 times

  **CODE\_STS** 2 years, 4 months ago

Got this in the exam today! Feb 28, 2023

upvoted 3 times

- 🗨️ 👤 **uffuchsi** 2 years, 4 months ago  
Received this in my exam today (22/02/2023). Selected copyIndex, copy, and dependsOn. Score 927.  
upvoted 10 times
- 🗨️ 👤 **saravanasanthosh** 2 years, 6 months ago  
Got this on exam 12/30/2022  
upvoted 2 times
- 🗨️ 👤 **OPT\_001122** 2 years, 6 months ago  
Thanks to all who have mentioned the exam date  
upvoted 2 times
- 🗨️ 👤 **[Removed]** 2 years, 9 months ago  
Do I need to remember these functions? The options are made in a way to confuse you and make a mistake. These are things developers just lookup in documentation. Not writing an ARM template every day  
upvoted 8 times
- 🗨️ 👤 **Rini100** 3 years, 2 months ago  
Got this on 20 Apr 2022.. (copyIndex, copy, dependsOn)  
upvoted 5 times
- 🗨️ 👤 **AZ204Cert** 3 years, 2 months ago  
Got this on 04/05/22 (selected copyIndex, copy, dependsOn)  
upvoted 4 times
- 🗨️ 👤 **petitbilly** 3 years, 3 months ago  
Got it in exam 03/22  
upvoted 3 times
- 🗨️ 👤 **AidenYoukhana** 3 years, 6 months ago  
CORRECT ANSWER.  
upvoted 2 times
- 🗨️ 👤 **Rev1201** 3 years, 9 months ago  
Answer is Correct!  
upvoted 3 times
- 🗨️ 👤 **glam** 4 years, 1 month ago  
Answer: copyIndex, copy, Dependson  
upvoted 3 times

## HOTSPOT -

You are developing an Azure Function App by using Visual Studio. The app will process orders input by an Azure Web App. The web app places the order information into Azure Queue Storage.

You need to review the Azure Function App code shown below.

```
public static class OrderProcessor
{
    [FunctionName("ProcessOrders")]
    public static void ProcessOrders([QueueTrigger("incoming-orders")]CloudQueueMessage myQueueItem, [Table("Orders")]ICollector<Order> tableBindings, TraceWriter log)
    {
        log.Info($"Processing Order: {myQueueItem.Id}");
        log.Info($"Queue Insertion Time: {myQueueItem.InsertionTime}");
        log.Info($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
        tableBindings.Add(JsonConvert.DeserializeObject<Order>(myQueueItem.AsString));
    }
    [FunctionName("ProcessOrders-Poison")]
    public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]CloudQueueMessage myQueueItem, TraceWriter log)
    {
        log.Error($"Failed to process order: {myQueueItem.AsString}");
        ...
    }
}
```

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Yes No

The code will log the time that the order was processed from the queue.

☐ ☐

When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.

☐ ☐

When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.

☐ ☐

The ProcessOrders function will output the order to an Orders table in Azure Table Storage.

☐ ☐

Suggested Answer:

### Answer Area

Yes No

The code will log the time that the order was processed from the queue.

☐ ☒

When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.

☒ ☐

When there are multiple orders in the queue, a batch of orders will be retrieved from the queue and the ProcessOrders function will run multiple instances concurrently to process the orders.

☒ ☐

The ProcessOrders function will output the order to an Orders table in Azure Table Storage.

☒ ☐

Box 1: No -

ExpirationTime - The time that the message expires.

InsertionTime - The time that the message was added to the queue.

Box 2: Yes -

maxDequeueCount - The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes -

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them.

By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.



Box 4: Yes -

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue>

  **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: No

It logs the following:

- ExpirationTime - The time that the message expires.
- InsertionTime - The time that the message was added to the queue.

Box 2: Yes

maxDequeueCount: The number of times to try processing a message before moving it to the poison queue. Default value is 5.

Box 3: Yes

When there are multiple queue messages waiting, the queue trigger retrieves a batch of messages and invokes function instances concurrently to process them. By default, the batch size is 16. When the number being processed gets down to 8, the runtime gets another batch and starts processing those messages. So the maximum number of concurrent messages being processed per function on one virtual machine (VM) is 24.



Box 4: Yes

`[Table("Orders")]ICollector<Order> table bindings`

And in the code it adds the order:



```
tableBindings.Add(JsonConvert.DeserializeObject<Object>(myQueueItem.AsString));
```

upvoted 144 times

  **OPT\_001122** 2 years, 7 months ago

very good and helpful explanation

upvoted 3 times

  **NombreFalso** 2 years, 4 months ago

Egg salad I mean excellent

upvoted 8 times

  **[Removed]** Highly Voted 2 years, 9 months ago

Microsoft thinks that a good developer should remember the default value for the dequeCount (and not forget that in a stress exam situation that there is that property, which is not shown in the code, very convenient).

upvoted 50 times

  **mr\_wilkrafal** 7 months, 4 weeks ago



Like, if not your comment, was gonna ask if we that image was updated or something...

upvoted 1 times

  **bgbgvf** Most Recent 1 year, 6 months ago



I think the answer is correct

upvoted 1 times

  **Anvsoc** 3 years, 6 months ago

correct

upvoted 1 times

  **tis\_truth** 3 years, 6 months ago

Would love to answer this correctly but the full code isn't showing. The maxDequeueCount value isn't showing in the question codeset provided. Does anyone have the full codeset?

upvoted 9 times

  **coffecold** 2 years, 9 months ago

maxDequeueCount; default =5

<https://docs.microsoft.com/fr-fr/azure/azure-functions/functions-bindings-storage-queue#hostjson-settings>

upvoted 3 times

  **[Removed]** 2 years, 9 months ago

It is not in the function but in the host.json. Microsoft is testing whether you forgot about that. I really do not know why they try to trick people but then it is Microsoft

upvoted 6 times



🗨️ 👤 **still6dark** 4 years, 1 month ago

Image isn't showing a full code? Please send me full code  
upvoted 8 times

🗨️ 👤 **j888** 3 years, 11 months ago

Same.. obviously something is missing  
upvoted 2 times

🗨️ 👤 **glam** 4 years, 1 month ago

Correct.  
upvoted 1 times

🗨️ 👤 **kapetan** 4 years, 2 months ago

The last statement is true: take a look at the input parameters:  
...[Table("Orders")]ICollection<Order> table bindings..  
and in the code it adds the order:  
`tableBindings.Add(JsonConvert.DeserializeObject<Object>(myQueueItem.AsString));`  
upvoted 2 times

🗨️ 👤 **idrisfl** 4 years, 3 months ago

Seems correct  
<https://docs.microsoft.com/fr-fr/azure/azure-functions/functions-bindings-storage-queue#hostjson-settings>  
upvoted 22 times

🗨️ 👤 **Marusyk** 4 years, 3 months ago

Answer is correct  
upvoted 3 times

🗨️ 👤 **Kuna\_Lambo** 4 years, 3 months ago

AZ-203 Topic 1 Q#9  
upvoted 2 times

## DRAG DROP -

You are developing a solution for a hospital to support the following use cases:

- ⇒ The most recent patient status details must be retrieved even if multiple users in different locations have updated the patient record.
- ⇒ Patient health monitoring data retrieved must be the current version or the prior version.
- ⇒ After a patient is discharged and all charges have been assessed, the patient billing record contains the final charges.

You provision a Cosmos DB NoSQL database and set the default consistency level for the database account to Strong. You set the value for Indexing Mode to Consistent.

You need to minimize latency and any impact to the availability of the solution. You must override the default consistency level at the query level to meet the required consistency guarantees for the scenarios.

Which consistency levels should you implement? To answer, drag the appropriate consistency levels to the correct requirements. Each consistency level may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

**Consistency levels**

Strong	Bounded Staleness
Consistent Prefix	Eventual

**Answer Area**

Return the most recent patient status.

Consistency level

Return health monitoring data that is no less than one version behind.

Consistency level

After patient is discharged and all charges are assessed, retrieve the correct billing data with the final charges.

Consistency level

**Suggested Answer:****Consistency levels**

Strong	Bounded Staleness
Consistent Prefix	Eventual

**Answer Area**

Return the most recent patient status.

Strong

Return health monitoring data that is no less than one version behind.

Bounded Staleness

After patient is discharged and all charges are assessed, retrieve the correct billing data with the final charges.

Eventual

**Box 1: Strong -**

Strong: Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

**Box 2: Bounded staleness -**

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is

"updates") of an item or by "t" time interval. When you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (t) by which the reads might lag behind the writes

**Box 3: Eventual -**


Eventual: There's no ordering guarantee for reads. In the absence of any further writes, the replicas eventually converge.

**Incorrect Answers:**


Consistent prefix: Updates that are returned contain some prefix of all the updates, with no gaps. Consistent prefix guarantees that reads never see out-of-order writes.

**Reference:**

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

 **AndresMza** Highly Voted 4 years, 3 months ago

Answers are correct  
upvoted 78 times

 **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: Strong

Box 2: Bounded staleness

Box 3: Eventual

Note: Consistent prefix: Updates that are returned contain some prefix of all the updates, with no gaps. Consistent prefix guarantees that reads never see out-of-order writes.

Reference:



<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

upvoted 34 times

  **TakumaK** 4 years ago

what is the Note for in your comment? you just copied it from the correct answer description which doesn't make sense to me.

upvoted 27 times

  **MrAzz** 1 year, 8 months ago

Consistent Prefix: This guarantees that a read operation will return the most recent version of the data written (the order is correct), but it could be outdated. It does not guarantee the most recent value of all writes anywhere.

For example, if data is written in A, B, C order, the user may get A, B or A, B, C. It will not get out of order writes such as A, C or A, C, B.

<https://www.mssqltips.com/sqlservertip/7158/azure-cosmos-db-consistency-levels-strong-bounded-staleness-session-consistent-prefix-eventual/>

upvoted 3 times

  **AhmedAbdelAziz** 1 year, 5 months ago

Consistent Prefix keep the order only within transaction , but with single write it behave like Eventual

upvoted 2 times


  **Thameur01**  3 months ago

Strong

Bounded Staleness

Strong: final charges = most recent data

upvoted 2 times

  **heptadecane** 7 months, 3 weeks ago

Strong, Strong, Eventual

> Return the most recent patient status : **\*\*Strong\*\***

> Strong consistency offers a linearizability guarantee. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

>

> Return health monitoring data that is no less than one version behind : **\*\*Strong\*\***

> Bounded staleness The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions of an item or by "t" time interval.

> However, For a single region account, the minimum value of K and T is 10 write operations or 5 seconds. For multi-region accounts the minimum value of K and T is 100,000 write operations or 300 seconds.

> Therefore, in order to ensure data is no less than once version behind, Strong consistency should be used.

> After patient is discharged and all charges are assessed, retrieve the correct billing data with the final charges : **\*\*Eventual\*\***

> In the absence of any further writes, the replicas eventually converge.

upvoted 1 times

  **Christian\_garcia\_martin** 10 months, 2 weeks ago

Box 3 is Eventual and not strong ,once all charges are assessed billing information wont change frequently and latency for retrieval is an important factor so .

Strong , Bounded staleness , Eventual

upvoted 1 times

  **DaniloPeres** 11 months, 3 weeks ago

Box 3 should be Strong

From ChatGPT

Key Consideration:

Correct Billing Data: This implies that the data must be accurate and reflect the final charges without any inconsistencies.

Decision:

Strong Consistency is the safest choice for retrieving the final billing data because it ensures that you get the latest and accurate charges immediately after all assessments are done.

Eventual Consistency might result in retrieving stale data if you read the billing information before the system has converged to the latest state. This could lead to inaccuracies in billing, which is unacceptable in a healthcare billing context.

upvoted 3 times

  **AhmedAbdelAziz** 1 year, 5 months ago

We all Agreed that

Box 1: Strong

Box 2: Bounded staleness


For Box 3 : it should be Strong as you should issue the bill with all the items added , for example in the Right Region you have item1 , item2, item3 if you use Bounded staleness or Eventual you may read in the time you issue the bill from read region item1, item2 only and for financial matter it's disaster and for the bill the order doesn't matter what matter is you should issue the bill with all items

upvoted 5 times

  **dddddd111** 1 year, 8 months ago



This is in my exam last Sept 20, 2023. I chose same answers.

upvoted 3 times

  **Ciupaz** 1 year, 7 months ago



Exam passed?

upvoted 1 times

  **agiler** 1 year, 9 months ago

for box 3, why should i risk to have a wrong bill because the database was not aligned yet?

upvoted 4 times

  **shekhar11** 1 year, 9 months ago



In my exam on 22sep 2023

upvoted 2 times

  **kingAzure** 1 year, 10 months ago



Got this Q on my exam 25/08/2023

upvoted 2 times

  **yawxir** 1 year, 9 months ago

Hi, How was your exam, considering the latest Open Book development ?



upvoted 1 times

  **jakobste** 1 year, 10 months ago

In a real world scenario, if the last charge is applied and then the final bill is immediately generated from a query against the database you might be in trouble with eventual consistency. However, if you can read the data as events in order asynchronous it would not be a problem.

I don't think i would go through the effort though, i would just use strong consistency to generate the final bill.

upvoted 2 times

  **yawxir** 1 year, 9 months ago

Agreed. I think the answer in last should be Strong consistency due to financially critical scenario.

upvoted 2 times

  **JH81** 2 years ago

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 4 times

  **Sleepers** 2 years ago



Got this question on 27/06/2023

upvoted 1 times

  **LanGo** 2 years, 2 months ago



Got this question today (17/04/2023). Answered strong, bounded staleness, eventual. Result: 846

upvoted 5 times

  **Papit** 2 years, 2 months ago

rcvd this question 04/2023

upvoted 1 times

  **Una20** 2 years, 3 months ago

received the qn in exam 02/2022

upvoted 2 times

## HOTSPOT -

You are configuring a development environment for your team. You deploy the latest Visual Studio image from the Azure Marketplace to your Azure subscription.

The development environment requires several software development kits (SDKs) and third-party components to support application development across the organization. You install and customize the deployed virtual machine (VM) for your development team. The customized VM must be saved to allow provisioning of a new team member development environment.

You need to save the customized VM for future provisioning.

Which tools or services should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Action	Tool or service
Generalize the VM.	<div><div>▼</div><div>Azure PowerShell</div><div>Visual Studio command prompt</div><div>Azure Migrate</div><div>Azure Backup</div></div>
Store images.	<div><div>▼</div><div>Azure Blob Storage</div><div>Azure Data Lake Storage</div><div>Azure File Storage</div><div>Azure Table Storage</div></div>

Suggested Answer:

**Answer Area**

Action	Tool or service
Generalize the VM.	<div><div>▼</div><div>Azure PowerShell</div><div>Visual Studio command prompt</div><div>Azure Migrate</div><div>Azure Backup</div></div>
Store images.	<div><div>▼</div><div>Azure Blob Storage</div><div>Azure Data Lake Storage</div><div>Azure File Storage</div><div>Azure Table Storage</div></div>

Box 1: Azure Powershell -

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks.

Before you begin, make sure that you have the latest version of the Azure PowerShell module.

You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image.

Box 2: Azure Blob Storage -

You can store images in Azure Blob Storage.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource#create-an-image-of-a-vm-using-powershell>

 **st003**  3 years, 8 months ago

The answer is right, it is show in AZ-900 as well.  
upvoted 58 times

 **clarionprogrammer** 3 years, 8 months ago

Powershell is correct.

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource>

upvoted 9 times

  **mlantonis** Highly Voted 3 years, 7 months ago

Box 1: Azure PowerShell

Creating an image directly from the VM ensures that the image includes all of the disks associated with the VM, including the OS disk and any data disks. Before you begin, make sure that you have the latest version of the Azure PowerShell module. You use Sysprep to generalize the virtual machine, then use Azure PowerShell to create the image.

Box 2: Azure Blob Storage

A VM Image is a collection of metadata and pointers to a set of VHDs (one VHD per disk) stored as page blobs in Azure Storage.

Reference:

<https://azure.microsoft.com/en-us/blog/vm-image-blog-post>

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/capture-image-resource>

upvoted 44 times

  **bgbgvf** Most Recent 1 year ago

I think the answer is right

upvoted 1 times

  **nlk0** 1 year, 5 months ago



Is it still relevant for az-204 since they removed VMs from scope in 28th of april?

upvoted 6 times

  **macobuzi** 1 year, 3 months ago

I would say better safe than sorry because Microsoft tends to give a sh\*t about scope.

upvoted 8 times

  **tom112** 1 year, 7 months ago

It's Visual Studio Command Prompt. You can't generalize VM to make an image in azure.

<Windows>

Sysprep removes all your personal account and security information, and then prepares the machine to be used as an image. For information about Sysprep, see Sysprep overview.

Once Sysprep has finished, set the status of the virtual machine to Generalized.

=> So, you generalize VM in the OS side(windows/linux/etc), and just set the status in Azure powershell.

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

upvoted 1 times



  **vizay** 1 week ago

To generalize a VM for imaging, you must:

Run Sysprep inside the VM

Then use Set-AzVM -Generalized – which is only available in Azure PowerShell

upvoted 1 times

  **IrshadF** 1 year, 9 months ago



It was asked in exam on 26/03/23 as well.

upvoted 1 times

  **sarmaria** 1 year, 9 months ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 4 times

  **Esward** 1 year, 11 months ago

Given answers are correct!

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

upvoted 2 times

🗨️ 👤 **Magis** 2 years, 1 month ago

I think that answer is incorrect.

Correct answer:

-Visual Studio Command Prompt as first of all you need to Generalize VM. Generalize means that you need to run SysPrep and you can do it from OS side only. So Azure PowerShell cannot be used at this point but you can do it from Visual Studio Command Prompt as it has integration with OS CMD.

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

- Second option is correct. Azure Blob Storage.

upvoted 4 times

🗨️ 👤 **AllNickNamesTaken** 1 year, 10 months ago

I agree. Everyone else appears to be answering the question "What tool should you use to create an image?" which would be PowerShell. But, that is not the question being asked. The question is "What tool should you use to generalize the image?" and you are correct with VC Cmd Prompt.

upvoted 1 times

🗨️ 👤 **rol204** 2 years, 2 months ago

why blob image I can not understand

upvoted 1 times

🗨️ 👤 **Dani\_ac7** 2 years, 4 months ago

Answer is correct, Powershell because we need arm Blob Storage templates to store virtual machine images, specifically Page Blobs

upvoted 1 times

🗨️ 👤 **Eltooth** 2 years, 6 months ago

Powershell

Blob Storage

upvoted 2 times

🗨️ 👤 **meoukg** 2 years, 9 months ago

Got it on 03/2022, I chose as Generalize the VM = Azure PowerShell ; Store images = Azure Blob Storage

upvoted 4 times

🗨️ 👤 **elequiel** 2 years, 1 month ago

thanks for coment date

upvoted 1 times

🗨️ 👤 **BobCui** 3 years ago

Why not Azure Backup for the first box?

upvoted 5 times

🗨️ 👤 **RochaG2** 2 years, 4 months ago

I've got no ideia.. and no one seems to explain it

upvoted 2 times

🗨️ 👤 **ozbonny** 3 years, 4 months ago

powershell and blob storage

upvoted 2 times

🗨️ 👤 **glam** 3 years, 7 months ago

correct.

upvoted 3 times

🗨️ 👤 **titombo** 3 years, 9 months ago

What is the problem to store the VM images on Azure File Storage?

upvoted 4 times

🗨️ 👤 **vizay** 1 week ago

Meant for file shares, not virtual disks

upvoted 1 times

🗨️ 👤 **DanVe** 3 years, 9 months ago



Page blobs are the standard storage area for VM  
upvoted 5 times

You are preparing to deploy a website to an Azure Web App from a GitHub repository. The website includes static content generated by a script. You plan to use the Azure Web App continuous deployment feature. You need to run the static generation script before the website starts serving traffic. What are two possible ways to achieve this goal? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Add the path to the static content generation tool to WEBSITE\_RUN\_FROM\_PACKAGE setting in the host.json file.
- B. Add a PreBuild target in the websites csproj project file that runs the static content generation script.
- C. Create a file named run.cmd in the folder /run that calls a script which generates the static content and deploys the website.
- D. Create a file named .deployment in the root of the repository that calls a script which generates the static content and deploys the website.

#### Suggested Answer: AD

A: In Azure, you can run your functions directly from a deployment package file in your function app. The other option is to deploy your files in the d:\home\site

\wwwroot directory of your function app (see A above).

To enable your function app to run from a package, you just add a WEBSITE\_RUN\_FROM\_PACKAGE setting to your function app settings.

Note: The host.json metadata file contains global configuration options that affect all functions for a function app.

D: To customize your deployment, include a .deployment file in the repository root.

You just need to add a file to the root of your repository with the name .deployment and the content:

[config]


command = YOUR COMMAND TO RUN FOR DEPLOYMENT

this command can be just running a script (batch file) that has all that is required for your deployment, like copying files from the repository to the web root directory for example.

Reference:

<https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script> <https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package>

Community vote distribution



minsma Highly Voted 4 years, 4 months ago

I think it is B and D

upvoted 76 times

rdemontis 4 years, 3 months ago

I think you're correct: in fact

- run.cmd is used only to start a project as a dll file <https://www.sohaibtariq.com/console-webjob/%C3%ACndex/>

- WEBSITE\_RUN\_FROM\_PACKAGE doesn't consent the execution of any script. You can only run your web project from a package (.zip file typically)

<https://docs.microsoft.com/bs-latn-ba/azure/azure-functions/run-functions-from-deployment-package>

<https://github.com/Azure/app-service-announcements/issues/84>

- Instead in .csproj file PreBuild Event Target you can specify any command to execute before the compilation and the application execution.

<https://docs.microsoft.com/en-us/visualstudio/msbuild/msbuild-targets?view=vs-2019>

<https://stackoverflow.com/questions/44818730/is-there-a-net-core-cli-pre-before-build-task>

<https://stackoverflow.com/questions/28916414/visual-studio-add-pre-build-event-that-always-runs-c-project>

upvoted 11 times

Vegetta95 3 years, 3 months ago

If the solution is built on the local machine, you can add target before build and do operations like copy files, run .bat(which can generate script content). However, if the repository is hosted on GitHub, any target will not work, because GitHub does not build apps

upvoted 2 times

Arrrqqq 3 years, 11 months ago

More direct info <https://docs.microsoft.com/en-us/azure/app-service/deploy-best-practices#net>

"By default, Kudu executes the build steps for your .NET application (dotnet build)" - so pre-build step should work there. There is similar note

for node.js

upvoted 3 times

🗨️  **[Removed]** 2 years, 9 months ago


Me: Do I need to know .Net shit to pass AZ-204?

Microsoft: No, it is only concept you need to know. It does not check your knowledge of .Net (or whatever it is)

Questions: .csproj file PreBuild Event Target


Me: Microsoft oh Microsoft.

upvoted 37 times

🗨️  **dimsok** 2 years, 4 months ago

This is not .NET shit

upvoted 6 times

🗨️  **ZodiaC** 3 years, 11 months ago

1000000% CORRECT !!!!!

upvoted 2 times

🗨️  **rolandcha** Highly Voted 4 years, 4 months ago

the answer is :


C,D

upvoted 31 times

🗨️  **Drgn** 4 years, 4 months ago

I agree C and D. The explanation of the Answers point to a Azure Function App, either the question was entered wrong (as they meant Azure Function App) into ExamTopics or the answers they chose were wrong.

upvoted 8 times

🗨️  **ACCP1** 4 years, 4 months ago

It could be A. It is used here for a web app. <https://docs.microsoft.com/en-us/azure/app-service/deploy-run-package>

upvoted 5 times

🗨️  **Robert12345Robert** 3 years, 12 months ago

Why is this upvoted? When you check the link you can read yourself that this will not run any scripts for you.

upvoted 10 times

🗨️  **[Removed]** 2 years, 9 months ago

But the repository could have a script and that will be run automatically by kudu ... Sorry, I do not know it myself whether kudu will be applied when it is from github. There are a lot of permutation combination and ways in which something can be done. To expect someone to remember all the unfortunate choices Microsoft made in these choices is ridiculous. But then it is my mistake I chose to do this certification from Microsoft

upvoted 1 times

🗨️  **Kitkit** 4 years, 4 months ago


Can you explain why you think c is the answer?

upvoted 3 times

🗨️  **clarionprogrammer** 4 years, 2 months ago

'C' makes no sense. Nothing exists on the Azure website in regard to a "run.cmd".

upvoted 6 times

🗨️  **dreamcoder** 3 years, 6 months ago

C, D is correct.

To customize your deployment, include a .deployment file in the repository root. For more information, see [Customize deployments and Custom deployment script](#).

<https://docs.microsoft.com/en-us/azure/app-service/deploy-continuous-deployment?tabs=github>

<https://github.com/projectkudu/kudu/wiki/Custom-Deployment-Script>

upvoted 3 times

🗨️  **Myk246** Most Recent 4 months, 2 weeks ago

Selected Answer: BD

To achieve the goal of running the static generation script before the website starts serving traffic, you can use the following two methods:

B. Add a PreBuild target in the website's csproj project file that runs the static content generation script.

This ensures that the script runs during the build process before the website is deployed.

D. Create a file named `.deployment` in the root of the repository that calls a script which generates the static content and deploys the website.

This file can specify custom deployment commands, including running your static content generation script before the website is deployed.

upvoted 1 times

🗳️ 👤 **DD7826** 5 months, 3 weeks ago

**Selected Answer: BD**

correct

upvoted 1 times

🗳️ 👤 **Faolba** 9 months, 2 weeks ago

**Selected Answer: CD**

ChatGPT says C and D

upvoted 1 times

🗳️ 👤 **HariB1992** 1 year, 2 months ago

**Selected Answer: BD**

Copilot Says B and D

upvoted 1 times

🗳️ 👤 **Isoldhe** 1 year, 4 months ago

**Selected Answer: BD**

Copilot AI says B and D:

Option B is correct because you can use the `PreBuild` target in the `csproj` file to execute a custom command or script before the project is built. This way, you can run the static content generation script and include the generated files in the project output.

Option D is correct because you can use the `.deployment` file in the root of the repository to customize the deployment process and specify a custom deployment script. This way, you can run the static content generation script and deploy the website using the custom script.

upvoted 3 times

🗳️ 👤 **Jass1nonly** 1 year, 4 months ago

A could NOT be the solution, since `WEBSITE_RUN_FROM_PACKAGE` is strictly used for running an app from a zip package. It doesn't involve any pre-build script support. It has to be B & D IMO.

upvoted 1 times

🗳️ 👤 **manopeydaikon** 1 year, 5 months ago

The correct answers are C and D.

C. Creating a file named `run.cmd` in the folder `/run` that calls a script which generates the static content and deploys the website is a valid approach. You can use this file to run any pre-processing or setup steps before the website starts serving traffic.

D. Creating a file named `.deployment` in the root of the repository that calls a script which generates the static content and deploys the website is also a valid approach. The `.deployment` file can be used to specify custom deployment steps, including running scripts before the deployment process.

Both options allow you to execute custom scripts or commands before the website starts serving traffic, making them suitable for running the static content generation script.

upvoted 1 times

🗳️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

Had cases:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 4 times

🗳️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

FYI:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 7 times

🗨️ 👤 **notedo** 2 years, 2 months ago

Thanks for commenting

upvoted 1 times

🗨️ 👤 **Saluk\_DE** 2 years, 3 months ago

Question was in Exam 2023-03-30

upvoted 2 times

🗨️ 👤 **PP2015** 2 years, 3 months ago

Did my exam on 3/29/2023. This question was on it. I Selected B and D. Score 850

upvoted 3 times

🗨️ 👤 **adilkhan** 2 years, 3 months ago

C D chatGPT

upvoted 1 times

🗨️ 👤 **DonH** 2 years ago

without any additional comment this answer is useless. Just typing it in ChatGPT doesn't mean it's helpful.

upvoted 6 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

ChatGPT is not always right, actually

upvoted 3 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

**Selected Answer: CD**

It has to be C, or D if you want to customize the deployment process, for example you want to run your tests before deploying (or after) and cancel the deployment if they fail?

Any of these commands will generate the files required to deploy your site, mainly:

.deployment - Contains the command to run for deploying your site.

deploy.cmd - Contains the deployment script (or deploy.sh if running on Mac/Linux)

upvoted 2 times

🗨️ 👤 **NombreFalso** 2 years, 4 months ago

**Selected Answer: BD**

A is wrong, it's never mentioned to be an Azure Function. There is nothing wrong with B and D.

upvoted 1 times

## DRAG DROP -

You are developing an application to use Azure Blob storage. You have configured Azure Blob storage to include change feeds.

A copy of your storage account must be created in another region. Data must be copied from the current storage account to the new storage account directly between the storage servers.

You need to create a copy of the storage account in another region and copy the data.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Select and Place:

**Actions****Answer Area**

Use AZCopy to copy the data to the new storage account.

Deploy the template to create a new storage account in the target region.

Export a Resource Manager template.

Create a new template deployment.

Modify the template by changing the storage account name and region.

**Suggested Answer:****Actions****Answer Area**

Use AZCopy to copy the data to the new storage account.

Deploy the template to create a new storage account in the target region.

Export a Resource Manager template.

Create a new template deployment.

Modify the template by changing the storage account name and region.

Create a new template deployment.

Export a Resource Manager template.

Modify the template by changing the storage account name and region.

Deploy the template to create a new storage account in the target region.

Use AZCopy to copy the data to the new storage account.



To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice.

The steps are:

- ⇒ Export a template.
- ⇒ Modify the template by adding the target region and storage account name.
- ⇒ Deploy the template to create the new storage account.
- ⇒ Configure the new storage account.
- ⇒ Move data to the new storage account.
- ⇒ Delete the resources in the source region.

Note: You must enable the change feed on your storage account to begin capturing and recording changes. You can enable and disable changes by using Azure

Resource Manager templates on Portal or Powershell.

Reference:

🗲️ 👤 **MariusN** Highly Voted 👍 4 years, 7 months ago

I think you first have to export the Resource Manager template before you can create a new template deployment. So, swap options 1 and 2 from the solution. In the first link of the solution's text, exporting is also considered as the first step.

upvoted 157 times

🗲️ 👤 **Kobee** 4 years, 7 months ago

Solution is right.

The option "Create a new Template deployment" is not present on the first link.

And you can first create the template if you want to export it. So Create first and then export the new template

upvoted 1 times

🗲️ 👤 **Kobee** 4 years, 7 months ago

Solution is right.

The option "Create a new Template deployment" is not present on the first link.

And you have to first create the template if you want to export it. So Create first and then Export the new template

upvoted 1 times

🗲️ 👤 **dancsita** 4 years, 4 months ago

you are exporting a different (existing) template, which will be used in a brand new template deployment

upvoted 1 times

🗲️ 👤 **pablocg** 4 years, 7 months ago

I agree with you. In the first link, in the step to modify the template it first creates the template deployment from the marketplace.

You can see the general steps to deploy an ARM template from the portal in the docs:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/quickstart-create-templates-use-the-portal#edit-and-deploy-the-template>

upvoted 2 times

🗲️ 👤 **bharatdilse** 4 years ago

AzCopy can be called from within the template also. That makes the deployment fully automatic

upvoted 2 times

🗲️ 👤 **victor** Highly Voted 👍 4 years, 5 months ago

Answer is Wrong. correct Sequence is

Export

Create

Modify

Deploy

AZ copy

upvoted 118 times

🗲️ 👤 **Robert12345Robert** 3 years, 12 months ago

I agree, see:

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

upvoted 4 times

🗲️ 👤 **noros5** 3 years, 4 months ago

Yes, confirmation: <https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

upvoted 6 times

🗲️ 👤 **msavx** Most Recent 🕒 4 months, 1 week ago

Someone has clearly messed with the screen shot... see the left column borders for the second and third options. The values have been shuffled.

According to <https://learn.microsoft.com/en-us/azure/operational-excellence/relocation-storage-account?tabs=azure-portal>. the correct order would be...

Export

Modify

Create

Deploy

AzCopy

upvoted 2 times

🗨️ 👤 **Christian\_garcia\_martin** 11 months ago

To all of yours dont know what is first Create or Modify.

Export

Modify

Create

Deploy

AZ copy.

you need to modify first because it would attempt to create a new storage account with the same name and in the same region, which is not allowed as storage account names must be unique across Azure.

upvoted 2 times

🗨️ 👤 **fuji36850** 1 year, 1 month ago

I think there are 2 ways:

Export > Modify(in notepad) > Create Deployment

Export > Create Deployment -> Modify (inside deployment window)

upvoted 5 times

🗨️ 👤 **james2033** 1 year, 1 month ago

1, export

2, create

3, modify

4, deploy

5, use AzCopy

upvoted 3 times

🗨️ 👤 **FeriAZ** 1 year, 4 months ago

1.Export a Resource Manager template.

2.Modify the template by changing the storage account name and region.

3.Create a new template deployment.

4.Use AZCopy to copy the data to the new storage account.

Before deploying the template to create a new storage account, you need to modify the exported template. Since you cannot have two storage accounts with the same name in Azure, you'll need to change the storage account name. You also need to update the region to reflect the target region where you want the new storage account to be created.

upvoted 2 times

🗨️ 👤 **AhmedAbdelAziz** 1 year, 4 months ago

Export

Modify

Create

Deploy

AZ copy

upvoted 7 times

🗨️ 👤 **11\_NickName\_11** 1 year, 7 months ago

What you all think about this:

Deploy the template to create a new storage account in the target region(Move section)?

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal>

upvoted 1 times

🗨️ 👤 **11\_NickName\_11** 1 year, 7 months ago

<https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal#move>

upvoted 1 times

🗨️ 👤 **11\_NickName\_11** 1 year, 7 months ago



My mistake, clearly states: "create a new storage account",so as far I'm agree with this one:

Export -



Create New Template Deployment

Modify

Deploy

AzCopy

upvoted 1 times

  **11\_NickName\_11** 1 year, 7 months ago

Sorry, I mean this one:

\* Export a template.


\* Modify the template by adding the target region and storage account name.

\* Create a new Template deployment

\* Deploy the template to create the new storage account.

\* Use AzCopy

upvoted 4 times

  **macobuzi** 1 year, 10 months ago

According to Microsoft document:[https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?](https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal)

[toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal](https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal)

- Export

- Modify

- Deploy

- Create

-AzCopy

upvoted 1 times

  **macobuzi** 1 year, 9 months ago

Sorry, correct order is:

- Export

- Modify



- Create

- Deploy

- AzCopy

Need to create the template before deploying it

upvoted 3 times

  **maqsoodshah** 2 years, 3 months ago

1. Export a template.
2. Modify the template by adding the target region and storage account name.
3. Deploy the template to create the new storage account.
4. Create a new template deployment
5. Use AZCopy to Copy data to the new storage account.

<https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

upvoted 3 times

  **macobuzi** 1 year, 10 months ago

You messed up at 3 and 4. You need to create the template before deploying it.

upvoted 3 times

  **Defa** 2 years, 3 months ago

Modify before create.

<https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

upvoted 2 times

  **Strawberry\_123** 2 years, 4 months ago

As per my understanding, correct answer is-Export/Create/Modify/Deploy/Az Copy

1. Export the template and download/save it.
2. Create a new template and load the previously exported template's json file.
3. Modify the required details.

4. Deploy it.

4. Use Az copy

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&tabs=azure-portal>  
upvoted 3 times

🗨️ 👤 **narenazure** 2 years, 5 months ago

<https://learn.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

Export

Create

Modify

Deploy

AZ copy

upvoted 3 times

🗨️ 👤 **Knightie** 2 years, 10 months ago

Creating a Template vs Creating a Template Deployment are two different things.

Export Template. Modify the Template, Create the Deployment, Deployment the Template, Do Copy.

upvoted 6 times

🗨️ 👤 **kbsivaiyer** 2 years, 11 months ago

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

it is export, modify, create, deploy, azcopy

upvoted 4 times

🗨️ 👤 **Wengatz** 2 years, 10 months ago

Thanks for the link, but according to that, the order is export > create > modify > deploy > azcopy

upvoted 3 times

🗨️ 👤 **BasitJamil** 2 years, 11 months ago

Export

New Template Manager

Modify

Deploy

AzCopy

upvoted 3 times

## DRAG DROP -

You are preparing to deploy an Azure virtual machine (VM)-based application.

The VMs that run the application have the following requirements:

- ⇒ When a VM is provisioned the firewall must be automatically configured before it can access Azure resources.
- ⇒ Supporting services must be installed by using an Azure PowerShell script that is stored in Azure Storage.

You need to ensure that the requirements are met.

Which features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Features	Requirement	Feature
Run Command	Firewall configuration	
Serial console	Supporting services script	
Hybrid Runbook Worker		
Custom Script Extension		

## Suggested Answer:

Features	Requirement	Feature
	Firewall configuration	Run Command
Serial console	Supporting services script	Hybrid Runbook Worker
Custom Script Extension		

## Reference:

<https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker> <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command>


 **agueda** Highly Voted 4 years, 3 months ago

1. Run Command
2. Customer Script Extension


Some question: <https://www.examtopycs.com/discussions/microsoft/view/12062-exam-az-203-topic-1-question-11-discussion/>  
upvoted 106 times

 **titombo** 4 years, 3 months ago

On this link is showing the same as we have now, 1. Run Command and 2. Hybrid Runbook Worker  
upvoted 3 times

 **BrettusMaximus** 4 years, 1 month ago



Cant use run to configure firewall. The run command uses the VM OS but the VM does not have access at that point.  
upvoted 2 times

 **fesioche** 3 years, 10 months ago

The Run Command feature enables virtual machine and application management and troubleshooting using scripts, and is available even when the machine is not reachable, for example if the guest firewall doesn't have the RDP or SSH port open.

Ref: <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm#run-command>

upvoted 10 times

  **pmsiva** 4 years, 3 months ago

The Custom Script Extension downloads and executes scripts on Azure virtual machines. This extension is useful for post deployment configuration, software installation, or any other configuration or management tasks. Scripts can be downloaded from Azure storage or GitHub, or provided to the Azure portal at extension run time.

upvoted 5 times

  **samraw83** 3 years ago

since you need to access the storage account for the services script for the installation you need the hybrid automation.

upvoted 1 times

  **mlantonis**  4 years, 1 month ago

Box 1: Run Command

This capability is useful in all scenarios where you want to run a script within a VM. It's one of the only ways to troubleshoot and remediate a VM that doesn't have the RDP or SSH port open, because of improper network or administrative user configuration.

Box 2: Customer Script Extension

The Custom Script Extension downloads and executes scripts on Azure virtual machines. This extension is useful for post deployment configuration, software installation, or any other configuration or management tasks. Scripts can be downloaded from Azure storage or GitHub, or provided to the Azure portal at extension run time. The Custom Script Extension integrates with Azure Resource Manager templates, and can be run using the Azure CLI, PowerShell, Azure portal, or the Azure Virtual Machine REST API.

upvoted 75 times

  **mlantonis** 4 years, 1 month ago

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows>

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm>

upvoted 10 times

  **gebpt**  1 year ago

1. Customer Script Extension

2. Run Command

Explanation (Chat GPT4o):

Custom Script Extension: Can be used to run scripts after the VM is started, but it is typically used for initial configuration. Best suited for configuring the firewall during the VM creation process before it gains access to Azure resources.

Run Command: Is more appropriate for running scripts that install supporting services on an already running VM. The script can be stored in Azure Storage and executed using 'Run Command'. Ideal for running scripts that install supporting services on an already running VM.

upvoted 1 times

  **manuelvigarcia** 1 year ago

According to referred documentation, Hybrid Runbook Worker is retiring, cannot be instantiated anymore, and is advised to migrate to Extension-based Worker.

upvoted 1 times

  **nlk0** 1 year, 11 months ago



Is it still relevant for az-204 since they removed VMs from scope from 28th april 2023

upvoted 2 times

  **macobuzi** 1 year, 10 months ago

As someone has stated, the question appeared on 29th May 2023, which means it is still within the scope. But seriously, I don't care whether it is out of scope or not, I just study everything from Examtopic, better safe than sorry!

upvoted 10 times

  **nlk0** 1 year, 11 months ago

is it still scope of az 204, since VMs were removed from scope in 28th april 2023?

upvoted 2 times

  **LuchoGonzalez** 2 years ago

- 1) Run command  
2) Customer script extension(<https://learn.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows>)  
upvoted 1 times

🗨️ 👤 **deathRac3** 2 years ago

Question was there for me on 29th May 2023  
upvoted 6 times

🗨️ 👤 **tom112** 2 years, 1 month ago

Azure Automation Agent-based User Hybrid Runbook Worker (Windows and Linux) will retire on 31 August 2024  
<https://learn.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker>  
upvoted 1 times

🗨️ 👤 **KarthikCS** 2 years, 2 months ago

Why are we tested on ridiculously low level details like this in these azure exams? It makes no sense. These AZ exams are the most senseless exams I've ever seen.  
upvoted 8 times

🗨️ 👤 **PP2015** 2 years, 3 months ago

Did my exam on 3/29/2023. This question was on it. Selected Run Command, Customer Script Extension. Score 850  
upvoted 5 times

🗨️ 👤 **Esward** 2 years, 5 months ago

Run command, Custom Script Extensions are correct!

<https://learn.microsoft.com/en-us/azure/virtual-machines/windows/run-command>  
upvoted 1 times

🗨️ 👤 **narenazure** 2 years, 5 months ago

1. Run
  2. CSE (The Custom Script Extension downloads and runs scripts on Azure virtual machines (VMs). This extension is useful for post-deployment configuration, software installation, or any other configuration or management task. You can download scripts from Azure Storage or GitHub, or provide them to the Azure portal at extension runtime.)
- upvoted 2 times

🗨️ 👤 **micro9000** 2 years, 7 months ago

It should be  
1. Run Command  
2. Custom Script Extension  
because there is no hybrid setup mentioned on the question  
upvoted 1 times

🗨️ 👤 **vcfvct** 2 years, 8 months ago

Got this on 10/21/2022. went with Run/CSE  
upvoted 3 times

🗨️ 👤 **coffecold** 2 years, 9 months ago

firewall configuration : hybrid runbook worker  
supporting services : Custom Script Extension

<https://learn.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm>

firewall configuration : must be executed on the VM and eventually on the host machine and must be done automatically.  
supporting services : once access to the VM, it can be run using an ARM  
upvoted 1 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

No  
Runbook worker  
- Requires manual installation  
- the firewall must be automatically configured  
upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

Run

CSE

upvoted 2 times

## HOTSPOT -

A company is developing a Node.js web app. The web app code is hosted in a GitHub repository located at <https://github.com/TailSpinToys/webapp>.

The web app must be reviewed before it is moved to production. You must deploy the initial code release to a deployment slot named review. You need to create the web app and deploy the code.

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
```

-Name myResourceGroup -Location \$location

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

-Name \$webappname -Location \$location -ResourceGroupName myResourceGroup -Tier Standard

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

-Name \$webappname -ResourceGroupName myResourceGroup -Slot review

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

```
$PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SourceSlotName review -DestinationSlotName production
```

## Suggested Answer:

## Answer Area

```
$gitrepo="https://github.com/TailSpinToys/webapp"
$webappname="TailSpinToysWeb"
$location="WestUS2"
```

-Name myResourceGroup -Location \$location

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

-Name \$webappname -Location \$location -ResourceGroupName myResourceGroup -Tier Standard

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

-Name \$webappname -ResourceGroupName myResourceGroup -Slot review

New-AzWebAppSlot  
New-AzWebApp  
New-AzAppServicePlan  
New-AzResourceGroup

```
$PropertiesObject = @{repoUrl = "$gitrepo";branch = "master";}
Set-AzResource -PropertyObject $PropertiesObject -ResourceGroupName myResourceGroup -ResourceType
Microsoft.Web/sites/slots/sourcecontrols -ResourceName $webappname/review/web -ApiVersion 2015-08-01 -Force
Switch-AzWebAppSlot -Name $webappname -ResourceGroupName myResourceGroup
-SourceSlotName review -DestinationSlotName production
```

Box 1: New-AzResourceGroup -

The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan -

The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location

Box 3: New-AzWebApp -

The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

Box 4: New-AzWebAppSlot -

The New-AzWebAppSlot cmdlet creates an Azure Web App slot.


Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroup?view=azps-2.3.2> <https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azappserviceplan?view=azps-2.3.2> <https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebapp?view=azps-2.3.2> <https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebappslot?view=azps-2.3.2>

 **AndresMza** Highly Voted 3 years, 9 months ago

Answer is correct

upvoted 87 times

 **mlantonis** Highly Voted 3 years, 7 months ago

Box 1: New-AzResourceGroup

The New-AzResourceGroup cmdlet creates an Azure resource group.

Box 2: New-AzAppServicePlan

The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location


Box 3: New-AzWebApp

The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

Box 4: New-AzWebAppSlot

The New-AzWebAppSlot cmdlet creates an Azure Web App slot.

upvoted 38 times

 **mlantonis** 3 years, 7 months ago

Reference:

<https://docs.microsoft.com/en-us/powershell/module/az.resources/new-azresourcegroup?view=azps-2.3.2>

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azappserviceplan?view=azps-2.3.2>

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebapp?view=azps-2.3.2>

<https://docs.microsoft.com/en-us/powershell/module/az.websites/new-azwebappslot?view=azps-2.3.2>

upvoted 8 times

 **bgbgvf** Most Recent 1 year ago

Answer is correct

upvoted 2 times

 **shekhar11** 1 year, 3 months ago


Question in my exam on 22sept 2023

upvoted 3 times

 **Saluk\_DE** 1 year, 9 months ago

Question was in Exam 2023-03-30

upvoted 4 times

 **sarmaria** 1 year, 9 months ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 2 times

 **Ayman99** 1 year, 10 months ago

Was on the exam 27.02.2023

upvoted 2 times

 **70PineApple** 1 year, 10 months ago



Got this in exam today..20/02/23

score: 817

upvoted 1 times

🗉 👤 **NombreFalso** 1 year, 10 months ago

As vanilla as it gets

upvoted 1 times

🗉 👤 **menlapwa** 1 year, 11 months ago

The answer make sense to me

upvoted 1 times

🗉 👤 **saravanasanthosh** 2 years ago

Got this on exam 12/30/2022

upvoted 1 times

🗉 👤 **SoftSol** 2 years ago

Correct, got this in the exam today dec 13, 2022

upvoted 1 times

🗉 👤 **uditrm12** 2 years ago

examtopic questions are enough to pass?

upvoted 1 times

🗉 👤 **coffecold** 2 years, 3 months ago

Finally an easy to answer question, without bad wordings

upvoted 1 times

🗉 👤 **Eltooth** 2 years, 6 months ago

ResourceGroup

ServicePlan

WebApp

WebAppSlot

upvoted 2 times

🗉 👤 **petitbilly** 2 years, 9 months ago

Got it in exam 03/22

upvoted 4 times

🗉 👤 **debanjan10** 3 years, 2 months ago

Always:

Resource Group / Group -> App Service Plan -> Web App -> Web App Slot -> Web App Source

upvoted 11 times

🗉 👤 **[Removed]** 3 years, 6 months ago

There is another question similar to this one but with AzureCLI. But share the same structure.

upvoted 3 times

## HOTSPOT -

You are developing an application that needs access to an Azure virtual machine (VM).

The access lifecycle for the application must be associated with the VM service instance.

You need to enable managed identity for the VM.

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

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

```
$vm = Get-AzVM -ResourceGroupName "ContosoRG" -Name "ContosoVM"
```

```
Update-AzVM -ResourceGroupName "ContosoRG" -VM $vm
```

-AssignIdentity:	▼
-IdentityId:	

\$SystemAssigned	▼
\$UserAssigned	

### Suggested Answer:

## Answer Area

```
$vm = Get-AzVM -ResourceGroupName "ContosoRG" -Name "ContosoVM"
```

```
Update-AzVM -ResourceGroupName "ContosoRG" -VM $vm
```

-AssignIdentity:	▼
-IdentityId:	

\$SystemAssigned	▼
\$UserAssigned	

Box 1: -IdentityType -

Enable system-assigned managed identity on an existing Azure VM:

To enable a system-assigned managed identity, use the -IdentityType switch on the Update-AzVM cmdlet (see below).

Box 2: \$SystemAssigned -

```
$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM
```

```
Update-AzVM -ResourceGroupName myResourceGroup -VM $vm -IdentityType SystemAssigned
```

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm>

 **agueda** Highly Voted 4 years, 3 months ago

The parameter should be "IdentityType", not "IdentityId" as it's stated in the reference link.

"SystemAssigned" is correct


upvoted 82 times

 **d0bermannn** 3 years, 10 months ago

as we see here <https://docs.microsoft.com/en-us/powershell/module/az.compute/update-azvm?view=azps-6.2.1>

there are both IdentityType and IdentityId parameters for update-azvm

upvoted 3 times

 **zolani** 1 year, 2 months ago

```
# Enable managed identity for the VM
```

```
$vmName = "MyVM"
```

```
$resourceGroupName = "MyResourceGroup"
```

```
# Set the VM identity type to SystemAssigned
```

```
Set-AzVM -ResourceGroupName $resourceGroupName -VMName $vmName -AssignIdentity $SystemAssigned
```

```
# Print a success message
```

```
Write-Host "Managed identity enabled for VM $vmName in resource group $resourceGroupName."
```

upvoted 1 times

🗳️ 👤 **TonyMel** 2 years, 3 months ago  
correct, in 2023Mar24, score: 904/1000  
upvoted 10 times

🗳️ 👤 **Shadoken** 3 years, 7 months ago  
I agree agueda. If you see the documentation its mandatory add "-IdentityType" tag.

<https://docs.microsoft.com/en-us/powershell/module/az.compute/update-azvm?view=azps-6.6.0#syntax>  
upvoted 2 times

🗳️ 👤 **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: -IdentityType

-IdentityType: The type of identity used for the virtual machine. Valid values are SystemAssigned, UserAssigned, SystemAssignedUserAssigned, and None.

-IdentityId: Specifies the list of user identities associated with the virtual machine. The user identity references will be ARM resource IDs in the form:

Box 2: \$SystemAssigned

There are two types of managed identities:

- System-assigned: Some Azure services allow you to enable a managed identity directly on a service instance. When you enable a system-assigned managed identity an identity is created in Azure AD that is tied to the lifecycle of that service instance. So when the resource is deleted, Azure automatically deletes the identity for you. By design, only that Azure resource can use this identity to request tokens from Azure AD.

- User-assigned: You may also create a managed identity as a standalone Azure resource. You can create a user-assigned managed identity and assign it to one or more instances of an Azure service. In the case of user-assigned managed identities, the identity is managed separately from the resources that use it.

upvoted 76 times

🗳️ 👤 **edengoforit** 3 years, 5 months ago  
This answer deserves a top comment  
upvoted 1 times

🗳️ 👤 **mlantonis** 4 years, 1 month ago  
Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm>

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types>

upvoted 6 times

🗳️ 👤 **Thameur01** Most Recent 3 months ago

IdentityType

SystemAssigned

upvoted 1 times

🗳️ 👤 **AryaNiaz** 10 months ago

Correct Answer.

\$vm = Get-AzVM -ResourceGroupName myResourceGroup -Name myVM

Update-AzVM -ResourceGroupName myResourceGroup -VM \$vm -IdentityType SystemAssigned

Reference: <https://learn.microsoft.com/en-us/entra/identity/managed-identities-azure-resources/how-to-configure-managed-identities?pivot=qs-configure-powershell-windows-vm#enable-system-assigned-managed-identity-on-an-existing-azure-vm>

upvoted 2 times

🗳️ 👤 **Christian\_garcia\_martin** 10 months, 2 weeks ago

first should be IdentityType , but if you don't have this choice which closely maps to -IdentityType is -AssignIdentity

upvoted 1 times

🗳️ 👤 **harbox** 1 year, 3 months ago

# Login to your Azure account

Connect-AzAccount

# Specify the resource group and VM name

\$resourceGroupName = "YourResourceGroup"

\$vmName = "YourVMName"

```
# Get the VM object
```

```
$vm = Get-AzVM -ResourceGroupName $resourceGroupName -Name $vmName
```

```
# Enable system-assigned managed identity
```

```
$vm = Set-AzVMIdentity -ResourceGroupName $resourceGroupName -VM $vm -AssignIdentity:$SystemAssigned
```

```
# Update the VM with the managed identity configuration
```


```
Update-AzVM -ResourceGroupName $resourceGroupName -VM $vm
```

upvoted 1 times

 **bgbgvfvf** 1 year, 6 months ago


Given answer is correct

upvoted 1 times

 **ENGs** 1 year, 8 months ago

On my exam 2023-10 before the Update of the Exam

upvoted 1 times

 **p2006** 1 year, 9 months ago

<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm#enable-system-assigned-managed-identity-on-an-existing-azure-vm>

upvoted 2 times

 **DonH** 2 years ago

Just for information: I just had this question on my AZ204 exam - 16-jun-2023.

I barely made it (with only 767 points) so I can't inform anyone if this answer is correct or not, just stating that this is an actual exam question.

upvoted 3 times

 **tom112** 2 years, 1 month ago

<IdentityType>

The type of identity used for the virtual machine. Valid values are SystemAssigned, UserAssigned, SystemAssignedUserAssigned, and None.

Type: Nullable<T>[ResourceIdentityType]

Accepted values: SystemAssigned, UserAssigned, SystemAssignedUserAssigned, None

upvoted 1 times

 **fkaracan** 2 years, 4 months ago

since \$ states they both are variables. given answer is correct.

upvoted 1 times

 **AllNickNamesTaken** 2 years, 4 months ago

Please correct me if I am wrong, but the "The access lifecycle for the application must be associated with the VM service instance" part looks like a red herring. Since it says that the \*access\* lifecycle must be associated with the service instance and not the \*managed identity lifecycle\*, then a user assigned identity would work here. Delete the VM and the application has no more access to it. If true, that would make the "IdentityID" parameter correct in this case.

upvoted 1 times

 **70PineApple** 2 years, 4 months ago

Got this in exam today..20/02/23

score: 817

upvoted 3 times

 **Priya0703** 2 years, 4 months ago

Got this question in exam today on 20-02-2023

upvoted 2 times

 **carlosghosn** 2 years, 7 months ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

 **micro9000** 2 years, 7 months ago

There's no correct option, it should be "IdentityType"

<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/qs-configure-powershell-windows-vm#enable-system-assigned-managed-identity-on-an-existing-azure-vm>

assigned-managed-identity-on-an-existing-azure-vm

We should update this question and the answer

upvoted 1 times

  **warchoon** 2 years, 4 months ago

I think it's a trick. \$SystemAssigned is a variable for ID

upvoted 1 times

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure

Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: A

In the Consumption hosting plan, resources are added dynamically as required by your functions.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>

Community vote distribution

B (100%)

 **AndresMza** Highly Voted 3 years, 9 months ago

Answer should be "No". Consumption plan can take up to several minutes to trigger the function. See note from <https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-storage-blob-triggered-function>.

"When your function app runs in the default Consumption plan, there may be a delay of up to several minutes between the blob being added or updated and the function being triggered. If you need low latency in your blob triggered functions, consider running your function app in an App Service plan."

upvoted 149 times

 **forgetfulalligator** 2 years ago

Yeap , you can also use a premium plan to use pre-warmed instances so no cold start


upvoted 3 times

 **rdemontis** 3 years, 9 months ago

In my opinion you're right. Often in a consumption plan the function starts immediately but it's not guaranteed and in production environment minimizing latency is very important as required in the question. For further info about cold starts i found an interesting article:

<https://azure.microsoft.com/en-us/blog/understanding-serverless-cold-start/#:~:text=In%20the%20context%20of%20Azure,haven't%20been%20called%20recently.>

upvoted 7 times

 **Esward** 1 year, 11 months ago

Use the app service plan and set always on property to true to start processing immediately

upvoted 2 times

 **[Removed]** Highly Voted 3 years, 9 months ago

Hi all.

I would say, that this has to be "No"...

Quote: "The process to produce a mobile-friendly version of the image must start in less than one minute."

And this is only a feature from the "Premium plan" (Perpetually warm instances to avoid any cold start)....

Consumption plan = up to several minutes...

upvoted 18 times

 **e3562ed** Most Recent 8 months ago

Selected Answer: B

for immediate start, it has to be app service plan

upvoted 1 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 1 month ago

**Selected Answer: B**

As rest of people says the consumption plan doesn't assure low latency

upvoted 1 times

🗨️ 👤 **Tarajee** 1 year, 3 months ago

Got this on 2023 sept

upvoted 2 times

🗨️ 👤 **Chris17** 1 year, 8 months ago

**Selected Answer: B**

The answer should be "No" because with consumption plan take more than 1 minute, in the before question. same to the answer, it is with app service plan.

upvoted 1 times

🗨️ 👤 **NombreFalso** 1 year, 10 months ago

The explanation is bull. Don't buy it. It's B.

upvoted 1 times

🗨️ 👤 **carlosghosn** 2 years, 1 month ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

🗨️ 👤 **dookebroyle** 2 years, 5 months ago

**Selected Answer: B**

Answer should be "no" after reviewing this article. The processing needs to begin sooner and Microsoft specifies the difference between an Event Grid trigger and a storage blob trigger.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=in-process%2Cextensionv5&pivots=programming-language-csharp>

"Minimizing latency: If your function app is on the Consumption plan, there can be up to a 10-minute delay in processing new blobs if a function app has gone idle. To avoid this latency, you can switch to an App Service plan with Always On enabled. You can also use an Event Grid trigger with your Blob storage account. For an example, see the Event Grid tutorial."

upvoted 5 times

🗨️ 👤 **Satish\_Babu** 2 years, 6 months ago

**Selected Answer: B**

Correction answer is B (No)

Because - When your function app runs in the default Consumption plan, there may be a delay of up to several minutes between the blob being added or updated and the function being triggered. If you need low latency in your blob triggered functions, consider running your function app in an App Service plan.

upvoted 2 times

🗨️ 👤 **azure900practice** 2 years, 6 months ago

**Selected Answer: B**

Function app should be in App service plan with 'Always On' setting as true. Consumption plan have cold starts and latency.

upvoted 1 times

🗨️ 👤 **Eltooth** 2 years, 6 months ago

**Selected Answer: B**

B is correct answer.

Consumption plan needs to warm up (up to 10 minutes) before processing.

upvoted 3 times

🗨️ 👤 **Freidrich** 2 years, 10 months ago

**Selected Answer: B**

The correct answer is B.

upvoted 2 times

🗨️ 👤 **sozturk88** 2 years, 10 months ago

**Selected Answer: B**

10 min delay for Consuming

upvoted 2 times

  **philsboies** 2 years, 10 months ago

**Selected Answer: B**

I'd say no



upvoted 2 times

  **Zamaletto** 2 years, 10 months ago

**Selected Answer: B**

I would go for B

upvoted 2 times

  **MFahd** 2 years, 11 months ago

The answer is NO, as the consumption plan have low latency and It can take some time to add files to blob and trigger function

upvoted 1 times



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 develop and deploy an Azure App Service API app to a Windows-hosted deployment slot named Development. You create additional deployment slots named Testing and Production. You enable auto swap on the Production deployment slot.

You need to ensure that scripts run and resources are available before a swap operation occurs.

Solution: Update the app with a method named statuscheck to run the scripts. Update the app settings for the app. Set the WEBSITE\_SWAP\_WARMUP\_PING\_PATH and WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES with a path to the new method and appropriate response codes.

Does the solution meet the goal?

A. No

B. Yes

#### Suggested Answer: A

These are valid warm-up behavior options, but are not helpful in fixing swap problems.

Instead update the web.config file to include the applicationInitialization configuration element. Specify custom initialization actions to run the scripts.

Note: Some apps might require custom warm-up actions before the swap. The applicationInitialization configuration element in web.config lets you specify custom initialization actions. The swap operation waits for this custom warm-up to finish before swapping with the target slot.


Here's a sample web.config fragment.

```
<system.webServer>
<applicationInitialization>
<add initializationPage="/" hostname="[app hostname]" />
<add initializationPage="/Home/About" hostname="[app hostname]" />
</applicationInitialization>
</system.webServer>
```

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#troubleshoot-swaps>

Community vote distribution



**Carlous** Highly Voted 3 years, 9 months ago

Should be YES?

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots>

You can also customize the warm-up behavior with one or both of the following app settings:

WEBSITE\_SWAP\_WARMUP\_PING\_PATH: The path to ping to warm up your site. Add this app setting by specifying a custom path that begins with a slash as the value. An example is /statuscheck. The default value is /.

WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES: Valid HTTP response codes for the warm-up operation. Add this app setting with a comma-separated list of HTTP codes. An example is 200,202 . If the returned status code isn't in the list, the warmup and swap operations are stopped. By default, all response codes are valid.

WEBSITE\_WARMUP\_PATH: A relative path on the site that should be pinged whenever the site restarts (not only during slot swaps). Example values include /statuscheck or the root path, /.

upvoted 85 times

**Tom87** 3 years, 8 months ago

I agree.



For ASP.NET and ASP.NET Core developers, setting app settings in App Service are like setting them in <appSettings> in Web.config or appsettings.json, but the values in App Service override the ones in Web.config or appsettings.json.

<https://docs.microsoft.com/en-us/azure/app-service/configure-common#configure-app-settings>

upvoted 5 times

**warchoon** 1 year, 10 months ago


NO. Auto swap should be on the source  
upvoted 5 times

  **altafpatel1984** 3 years, 1 month ago


Question is not only for warm-up but also to execute custom script. So Answer No is correct.  
upvoted 11 times

  **mcanic** 3 years ago


it states that a new method called statuscheck is added to the application that executes these scripts. if you add the path to /statuscheck in WEBSITE\_SWAP\_WARMUP\_PING\_PATH the application executes the scripts  
upvoted 10 times

  **Santileo** 3 years, 8 months ago



This guy is right. Better url: <https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up>  
upvoted 6 times

  **Percy2112** 3 years, 6 months ago

I Agree.  
upvoted 3 times

  **simonverma**  3 years, 8 months ago

The answer is right to be marked as 'No' as the warm up is done correctly but the web config part is not mentioned i.e, AppInitialization.  
upvoted 23 times

  **lugospod** 2 years, 11 months ago

you dont need appinit tif you set WEBSITE\_SWAP\_WARMUP\_PING\_PATH and response... it is an alternate method to achieve the same thing.  
upvoted 6 times

  **HariB1992**  8 months, 1 week ago



Yes, the solution meets the goal!  
upvoted 3 times

  **Isoldhe** 10 months, 1 week ago



Copilot AI says No:

The solution does not meet the goal. Updating the app with a method named statuscheck to run the scripts and setting the WEBSITE\_SWAP\_WARMUP\_PING\_PATH and WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES app settings are not sufficient to ensure that scripts run and resources are available before a swap operation occurs. These settings are used to configure the health check for the swap operation, which determines whether the source slot is ready to be swapped with the target slot. However, the health check does not guarantee that the scripts run and resources are available, as it only verifies that the app responds to HTTP requests with the expected status codes.  
upvoted 3 times

  **Jass1nonly** 10 months, 2 weeks ago



Quote "You need to ensure that scripts run and resources are available before a swap operation occurs."

Note from Microsoft "The <applicationInitialization> configuration element is part of each app start-up, whereas the two warm-up behavior app settings apply only to slot swaps."

conclusion: for this case where swap is concerned, The proposed solution in the question with appsettings WEBSITE\_SWAP\_WARMUP\_PING\_PATH and WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES is correct.

Hence answer B. Yes.

upvoted 1 times

  **AhmedAbdelAziz** 10 months, 4 weeks ago



Should be YES

The <applicationInitialization> configuration element is part of each app start-up, whereas the two warm-up behavior app settings apply only to slot swaps.

WEBSITE\_SWAP\_WARMUP\_PING\_PATH: The path to ping over HTTP to warm up your site. Add this app setting by specifying a custom path that begins with a slash as the value. An example is /statuscheck. The default value is /.

WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES: Valid HTTP response codes for the warm-up operation. Add this app setting with a comma-separated list of HTTP codes. An example is 200,202 . If the returned status code isn't in the list, the warmup and swap operations are stopped. By default, all response codes are valid.

upvoted 1 times

🗨️ 👤 **Nihilist11** 11 months, 2 weeks ago

Jan-21-2024 - This was in my exam, went with YES - Score 740

Contosso case study [Couldnt find here]

upvoted 4 times

🗨️ 👤 **paraballa** 11 months ago

How many questions were there that came from ExamTopics?

upvoted 1 times

🗨️ 👤 **Tarajee** 1 year, 3 months ago

Got this on 2023sept

upvoted 1 times

🗨️ 👤 **mihailos** 1 year, 4 months ago

Got it in exam 28/08/23. Went with YES for this. Scored 912

upvoted 4 times

🗨️ 👤 **yawxir** 1 year, 3 months ago

Hey,

Glad you pass. Can you guide more on exam pattern as this should be open book now ?

upvoted 3 times

🗨️ 👤 **JH81** 1 year, 6 months ago

**Selected Answer: A**

Answer is NO. While this solution may work put on your MS thinking hat because modifications to the app should not be necessary when there are other ways to get this to work with no app modifications.

upvoted 1 times

🗨️ 👤 **red0400** 1 year, 9 months ago

**Selected Answer: B**

Answer is YES

upvoted 1 times

🗨️ 👤 **uffuchsi** 1 year, 10 months ago

Surely it's yes?

You can customize the warm-up behavior with one or both of the following app settings:

WEBSITE\_SWAP\_WARMUP\_PING\_PATH: The path to ping to warm up your site. Add this app setting by specifying a custom path that begins with a slash as the value. An example is /statuscheck. The default value is /.

WEBSITE\_SWAP\_WARMUP\_PING\_STATUSES: Valid HTTP response codes for the warm-up operation. Add this app setting with a comma-separated list of HTTP codes. An example is 200,202 . If the returned status code isn't in the list, the warmup and swap operations are stopped. By default, all response codes are valid.

WEBSITE\_WARMUP\_PATH: A relative path on the site that should be pinged whenever the site restarts (not only during slot swaps). Example values include /statuscheck or the root path, /.

<https://docs.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up>

upvoted 1 times

🗨️ 👤 **rgullini** 1 year, 12 months ago

**Selected Answer: B**

The solution doesn't include the scripts. Answer is NO.

upvoted 1 times

🗨️ 👤 **Lucifer14** 2 years ago

In 28-12-2022 exam

upvoted 1 times

🗨️ 👤 **memooo** 2 years ago

**Selected Answer: B**

Correct answer is B:

Reference:

<https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots#specify-custom-warm-up>

upvoted 1 times

  **mdg3501** 2 years ago

got this on 2022-12-7

upvoted 1 times

  **nvtienanh** 2 years ago

Got this question on 2-Dec-2022 exam. Went with most voted

upvoted 3 times

  **uditrm12** 2 years ago

examtopics questions are enough to pass?

upvoted 2 times

## HOTSPOT -

You create the following PowerShell script:

```
$source = New-AzScheduledQueryRuleSource -Query 'Heartbeat | where TimeGenerated > ago(1h)' -DataSourceId "contoso"
$schedule = New-AzScheduledQueryRuleSchedule -FrequencyInMinutes 60 -TimeWindowInMinutes 60
$triggerCondition = New-AzScheduledQueryRuleTriggerCondition -ThresholdOperator "LessThan" -Threshold 5
$aznsActionGroup = New-AzScheduledQueryRuleAznsActionGroup -ActionGroup "contoso" -EmailSubject "Custom email subject"
-CustomWebhookPayload "{ 'alert': '#alertrulename', 'IncludeSearchResults': true }"
$alertingAction = New-AzScheduledQueryRuleAlertingAction -AznsAction $aznsActionGroup -Severity "3" -Trigger $triggerCondition
New-AzScheduledQueryRule -ResourceGroupName "contoso" -Location "eastus" -Action $alertingAction -Enabled $true
-Description "Alert description" -Schedule $schedule -Source $source -Name "Alert Name"
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	<input type="radio"/>	<input type="radio"/>
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	<input type="radio"/>	<input type="radio"/>
The log alert is scheduled to run every two hours.	<input type="radio"/>	<input type="radio"/>

### Answer Area

Suggested Answer:

Statements	Yes	No
A log alert is created that sends an email when the CPU percentage is above 60 percent for five minutes.	<input type="radio"/>	<input checked="" type="radio"/>
A log alert is created that sends an email when the number of virtual machine heartbeats in the past hour is less than five.	<input checked="" type="radio"/>	<input type="radio"/>
The log alert is scheduled to run every two hours.	<input type="radio"/>	<input checked="" type="radio"/>

Box 1: No -

The AzScheduledQueryRuleSource is Heartbeat, not CPU.

Box 2: Yes -

The AzScheduledQueryRuleSource is Heartbeat!

Note: New-AzScheduledQueryRuleTriggerCondition creates an object of type Trigger Condition. This object is to be passed to the command that creates Alerting Action object.

Box 3: No -

The schedule is 60 minutes, not two hours.

-FrequencyInMinutes: The alert frequency.

-TimeWindowInMinutes: The alert time window

The New-AzScheduledQueryRuleSchedule command creates an object of type Schedule. This object is to be passed to the command that creates Log Alert Rule.

Reference:

🗲️ 👤 **mlantonis** Highly Voted 3 years, 7 months ago

I agree.

No

Yes

No

upvoted 92 times

🗲️ 👤 **TonyMel** 1 year, 9 months ago  
correct, in 2023Mar24, score: 904/1000  
upvoted 9 times

🗲️ 👤 **wolf\_lu** Highly Voted 3 years, 6 months ago

I agree.

NO

YES

NO

upvoted 10 times

🗲️ 👤 **barretowski** 2 years, 9 months ago  
don't confuse the staff, the answer is: no, yes, no  
upvoted 2 times

🗲️ 👤 **bgbgvf** Most Recent 1 year ago

Given answer is correct

upvoted 1 times

🗲️ 👤 **JH81** 1 year, 6 months ago  
Got this on 6/28/2023 and passed with 850. Went with No, Yes, No.  
upvoted 4 times

🗲️ 👤 **aragones** 1 year, 7 months ago  
Got this 2023-05-12.

FYI:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 1 times

🗲️ 👤 **ucskips** 1 year, 9 months ago

1. No

2. Yes

3. No

I used the last command to get the following understanding

My logic:

Commands will send an email

Where threshold is < 5

Every 60 mins

Where heartbeat is > 1h ago

upvoted 1 times

🗲️ 👤 **70PineApple** 1 year, 10 months ago

Got this in exam today..20/02/23

score: 817

upvoted 2 times

🗲️ 👤 **Priya0703** 1 year, 10 months ago

Got this question on exam today 20-02-2023.

upvoted 2 times

🗨️ 👤 **18Marks** 2 years ago

Got it in exam. Use highly voted answer

upvoted 2 times

🗨️ 👤 **uditrm12** 2 years ago

examtopics questions are enough to pass?

upvoted 1 times

🗨️ 👤 **70PineApple** 1 year, 10 months ago

yes they are

upvoted 1 times

🗨️ 👤 **70PineApple** 1 year, 10 months ago

yes they are

upvoted 2 times

🗨️ 👤 **Nokaido** 2 years, 3 months ago

What about the 60 min trigger frequency + the 60 min recording window? As far as I know it only triggers after the previous trigger is finished. It should be NO, YES, YES?

upvoted 2 times

🗨️ 👤 **JustHereToLearn** 2 years ago

No, the answer should be No, Yes, No as per this link <https://learn.microsoft.com/en-us/javascript/api/@azure/arm-monitor/schedule?view=azure-node-latest>

upvoted 3 times

🗨️ 👤 **nuiowmokka** 2 years, 4 months ago

Did anyone get this question in real exam?

upvoted 3 times

🗨️ 👤 **ImPavan** 2 years ago

I got today 28-12-2022

upvoted 1 times

🗨️ 👤 **SivajiTheBoss** 2 years, 9 months ago

Answer is correct: NO, YES, NO

upvoted 1 times

🗨️ 👤 **yaroo1** 3 years, 2 months ago

Answer is correct

upvoted 3 times

🗨️ 👤 **Percy2112** 3 years, 6 months ago

The Answers are Correct:

No

Yes

No

upvoted 3 times

🗨️ 👤 **Arul4** 3 years, 6 months ago

No

Yes

No

are the correct answers.

upvoted 4 times

## DRAG DROP -

You are developing an Azure Function app.

The app must meet the following requirements:

- ☞ Enable developers to write the functions by using the Rust language.
- ☞ Declaratively connect to an Azure Blob Storage account.

You need to implement the app.

Which Azure Function app features should you use? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

**Features**

Custom handler

Extension bundle

Trigger

Runtime

Policy

Hosting plan

**Answer Area****Requirement****Feature**

Enable developers to write the functions by using the Rust language.

Feature

Declaratively connect to an Azure Blob Storage account.

Feature

**Suggested Answer:****Features**

Extension bundle

Runtime

Policy

Hosting plan

**Answer Area****Requirement****Feature**

Enable developers to write the functions by using the Rust language.

Custom handler

Declaratively connect to an Azure Blob Storage account.

Trigger

Box 1: Custom handler -


Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: Trigger -

Functions are invoked by a trigger and can have exactly one. In addition to invoking the function, certain triggers also serve as bindings. You may also define multiple bindings in addition to the trigger. Bindings provide a declarative way to connect data to your code.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other> <https://docs.microsoft.com/en-us/dotnet/architecture/serverless/azure-functions>

 **abdou1987** Highly Voted 3 years, 6 months ago

Box 1: Custom handler

Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: extension bundles



is needed to support the bindings and triggers that you use

[https://docs.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers?WT.mc\\_id=thomasmaurer-blog-thmaure#bindings-support](https://docs.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers?WT.mc_id=thomasmaurer-blog-thmaure#bindings-support)  
upvoted 84 times

  **[Removed]** 1 year, 10 months ago

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers>

- With custom handlers, you can use triggers and input and output bindings via extension bundles.

So: Custom Handler and Extension bundles.

upvoted 3 times

  **Stel0Papad4** 7 months ago



It may be so but, ultimately you do need Triggers to do Box 2's action so, maybe the Trigger option is the correct answer for Box 2.

upvoted 2 times

  **Stel0Papad4** 7 months ago

(+in addition) It's like asking me --> To create a VM you need to use the portal or the bash and im telling you that no you cant use them. You need to have a subscription first. The correct answer is the subscription :P

upvoted 1 times

  **mlantonis** Highly Voted 3 years, 7 months ago

Answer is correct.

Box 1: Custom handler

Custom handlers can be used to create functions in any language or runtime by running an HTTP server process, for example Go or Rust.

Box 2: Trigger



Functions are invoked by a trigger and can have exactly one. In addition to invoking the function, certain triggers also serve as bindings. You may also define in addition to the trigger. Bindings provide a declarative way to connect data to your code.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other>

<https://docs.microsoft.com/en-us/dotnet/architecture/serverless/azure-functions>

upvoted 75 times

  **Esward** 1 year, 11 months ago

Agreed

upvoted 1 times

  **TakumaK** 3 years, 6 months ago

incorrect!

upvoted 2 times

  **stevenwu** 3 years, 6 months ago

What's your answer? You can just point out incorrect without any updates.

upvoted 18 times

  **azurelearner666** 3 years, 6 months ago

your incorrect is incorrect X2.

The answer is correct, Custom Handler for supporting additional languages in azure functions and a trigger for adding a binding to a function.

upvoted 11 times

  **aruni\_mishra** 2 years, 1 month ago

BOX 2: extension bundles

from MS:

With custom handlers, you can use triggers and input and output bindings via extension bundles.

[https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers?WT.mc\\_id=thomasmaurer-blog-thmaure#bindings-](https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers?WT.mc_id=thomasmaurer-blog-thmaure#bindings-support)

support~:text=With%20custom%20handlers%2C%20you%20can%20use%20triggers%20and%20input%20and%20output%20bindings%20via%20exte

upvoted 7 times

  **mmdex** 1 year, 10 months ago

Box 2: Extension Bundles

Yes, you could use a Trigger, but then you have only ONE trigger per function. The question says nothing about how the function is being triggered, only that it connects to Blob Storage. What if you wanted to trigger the function by HTTP request and then perform some action in Blob Storage? You can't do that with a trigger.

upvoted 8 times

🗳️ 👤 **neelkanths** Most Recent 8 months, 1 week ago

Box2: extension bundles -

As per question it is looking to connect declaratively with blob storage account, to connect declaratively one needs to use bindings - as per below article

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings?tabs=isolated-process%2Cpython-v2&pivots=programming-language-csharp>

therefore trigger can't be the answer.

upvoted 1 times

🗳️ 👤 **FeriAZ** 10 months, 3 weeks ago

Custom Handler.

Extension Bundle.

To "declaratively connect" to an Azure Blob Storage account refers to the ability to specify in the function's configuration how and to what resources the function connects, without needing imperative programming to establish the connection. While a trigger can start a function based on activity in a Blob Storage (like creating or adding a new blob), it's the bindings that are more directly involved in the declarative connection to data sources or sinks.

upvoted 3 times

🗳️ 👤 **AhmedAbdelAziz** 10 months, 4 weeks ago

Box 1: Custom handler

Box 2: trigger

extension bundles allow you to use trigger in the custom handler but it can't do the job alone you have to use trigger to fire the function

Custom handlers are best suited for situations where you want to:

- Implement a function app in a language that's not currently offered out-of-the box, such as Go or Rust.
- Implement a function app in a runtime that's not currently featured by default, such as Deno.

With custom handlers, you can use triggers and input and output bindings via extension bundles.

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-custom-handlers>

upvoted 1 times

🗳️ 👤 **shekhar11** 1 year, 3 months ago

Question in my exam on 22sept 2023

upvoted 3 times

🗳️ 👤 **Tarajee** 1 year, 3 months ago

Got this on 2023sept

upvoted 4 times

🗳️ 👤 **Kanasan** 1 year, 4 months ago

For the second I vote for 'Trigger':

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-triggers-bindings?tabs=isolated-process%2Cpython-v2&pivots=programming-language-csharp#add-bindings-to-a-function>

upvoted 1 times

🗳️ 👤 **kvtborad** 1 year, 4 months ago

I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city. All questions are from ExamTopics.

upvoted 1 times

🗳️ 👤 **RikinPatel** 1 year, 3 months ago

There are 2 most voted ans.. so can you please specify which most voted you have selected?

upvoted 5 times

🗳️ 👤 **Firo** 1 year, 5 months ago

Can anyone mention what the answer is for box2, there are 2 answers that got most voted which one is the correct one.

upvoted 2 times

🗳️ 👤 **JH81** 1 year, 6 months ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.

upvoted 1 times

🗨️ 👤 **tom112** 1 year, 7 months ago

For Box 2: Trigger. (But Binding would make more sense if there is.)

Extension bundles allow non-.NET functions apps to use the same bindings without having to deal with the .NET infrastructure.  
-> so it's not for general trigger or bindings setting.

Reference: Register Azure Functions binding extensions

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-register>

upvoted 1 times

🗨️ 👤 **aragones** 1 year, 7 months ago

Got this 2023-05-12.

FYI:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 1 times

🗨️ 👤 **adilkhan** 1 year, 9 months ago

Box 1: Custom handler

Box 2: extension bundles CHATGPT

upvoted 1 times

🗨️ 👤 **saravanasanthosh** 2 years ago

Got this on exam 12/30/2022

upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 1 month ago

Box 2: extension bundles or Triggers? which one is correct ? both answers are highly voted with valid comments

upvoted 2 times

🗨️ 👤 **guchao2000** 1 year, 11 months ago

Trigger

<https://learn.microsoft.com/en-us/azure/azure-functions/create-first-function-vs-code-other?tabs=rust%2Cwindows>

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 3 months ago

It has to be custom handler and extension bundle if the question has to be read together. But another of those questions that does not make much sense. People who want to create in RUST should use an app service instead of using a function app. It does not add a lot of value to play into these dark areas. Microsoft just adds features that are half baked instead of focusing on fixing the issues or making things consistent. Bad product management.

upvoted 5 times

## HOTSPOT -

You are developing an ASP.NET Core web application. You plan to deploy the application to Azure Web App for Containers.

The application needs to store runtime diagnostic data that must be persisted across application restarts. You have the following code:

```
public void SaveDiagData(string data)
{
    var path = Environment.GetEnvironmentVariable("DIAGDATA");
    File.WriteAllText(Path.Combine(path, "data"), data);
}
```

You need to configure the application settings so that diagnostic data is stored as required.

How should you configure the web app's settings? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

### App setting

▼
LOCALAPPDATA
WEBSITE_LOCALCACHE_ENABLED
DOTNET_HOSTING_OPTIMIZATION_CACHE
WEBSITES_ENABLE_APP_SERVICE_STORAGE
DIAGDATA

### Value

true
▼
/home
/local
D:\home
D:\local

Suggested Answer:

## Answer Area

### App setting

▼
LOCALAPPDATA
WEBSITE_LOCALCACHE_ENABLED
DOTNET_HOSTING_OPTIMIZATION_CACHE
WEBSITES_ENABLE_APP_SERVICE_STORAGE
DIAGDATA

### Value

true
▼
/home
/local
D:\home
D:\local

Box 1: If WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE

If WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE setting is unspecified or set to true, the /home/ directory will be shared across scale instances, and files written will persist across restarts

Box 2: /home -

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq>

 **zinza**  4 years, 5 months ago

I think the printed image is creating confusion. The correct answers (according to the image) are:

WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE=true

DIAGDATA=/home

upvoted 108 times

 **lugospod** 3 years, 5 months ago



How do you know it is Linux and not Windows?

upvoted 2 times

  **ScubaDiver123456** 3 years, 5 months ago

For Windows,it would be C:\Home, which is not a provided option.. but /home is provided.

<https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-windows>  
upvoted 15 times

  **Tej\_007** 9 months, 4 weeks ago

and also

By default, persistent storage is enabled on Windows custom containers. To disable it, set the WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE app setting value to false

upvoted 1 times

  **azurelearner666** 4 years ago



yep, correct. Same as the "official solution" :)

upvoted 6 times

  **GCMan**  4 years, 7 months ago

Given answer correct.

upvoted 25 times

  **Esward** 2 years, 5 months ago

Agreed

upvoted 1 times

  **Kanasan**  1 year, 10 months ago

The answer is correct, but explanation is not:

If WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE setting is unspecified or set to false, the /home/ directory will not be shared across scale instances, and files written will not persist across restarts.

<https://learn.microsoft.com/en-us/troubleshoot/azure/app-service/faqs-app-service-linux#i-m-using-my-own-custom-container--i-want-the-platform-to-mount-an-smb-share-to-the---home---directory->

upvoted 4 times

  **micro9000** 2 years, 7 months ago

The answer is correct

<https://learn.microsoft.com/en-us/azure/app-service/reference-app-settings?tabs=kudu%2Cdotnet#custom-containers>

upvoted 1 times

  **[Removed]** 2 years, 9 months ago

Custom container itself is a dark area of app service and expecting human beings to remember the property names and the possible values for these settings is very remarkable.

upvoted 10 times

  **[Removed]** 2 years, 9 months ago

I have to just remember all these properties? Are developers not allowed to use documentation in real life? Does search work in Redmond? Microsoft should really think what exactly they achieve from these random questions without a thought on what they want the developers to know or remember.

upvoted 15 times

  **Eltooth** 3 years ago

WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE

/home

upvoted 5 times



  **SivajiTheBoss** 3 years, 3 months ago

Answer is correct:

WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE=true


DIAGDATA=/home

upvoted 2 times

  **upadhyayavi** 3 years, 3 months ago

Can't D:/home and /home both be correct in this case??

upvoted 1 times

  **Ami\_Nou** 3 years, 1 month ago

For anyone confused, in Windows it is C:\Home

upvoted 5 times

🗨️ 👤 **upadhyayavi** 3 years, 3 months ago

Answer to both the options: <https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-windows#use-persistent-shared-storage>

upvoted 1 times

🗨️ 👤 **petitbilly** 3 years, 3 months ago

Here there is the specific documentation for this question: <https://docs.microsoft.com/en-us/azure/app-service/faq-app-service-linux#i-m-using-my-own-custom-container-i-want-the-platform-to-mount-an-smb-share-to-the---home---directory->

upvoted 1 times

🗨️ 👤 **qiw** 3 years, 5 months ago

WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE=true

By default, persistent storage is disabled on custom containers and the setting is exposed in the app settings. To enable it, set the WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE app setting value to true

<https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux>

DIAGDATA=/home

You can use the /home directory in your custom container file system to persist files across restarts and share them across instances. The /home directory is provided to enable your custom container to access persistent storage.

<https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux>

upvoted 5 times

🗨️ 👤 **mlantonis** 4 years, 1 month ago

Correct Answer:

If WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE setting is unspecified or set to false, the /home/ directory will not be shared across scale instances, and files written will not persist across restarts. Explicitly setting WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE to true will enable the mount.

Box 1: WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE

Box 2: /home

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/containers/app-service-linux-faq>

upvoted 16 times

🗨️ 👤 **glam** 4 years, 1 month ago

WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE=true

DIAGDATA=/home

upvoted 3 times

🗨️ 👤 **JoelnOregon** 4 years, 1 month ago

glam has the answers, y'all.

upvoted 3 times

🗨️ 👤 **Tom87** 4 years, 1 month ago

On Windows the directory would be C:\home. Since this is not mentioned in possible answers, we are on Linux environment and the answer is correct.

<https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-windows#use-persistent-shared-storage>

upvoted 4 times

🗨️ 👤 **glam** 4 years, 1 month ago

WEBSITES\_ENABLE\_APP\_SERVICE\_STORAGE=true

DIAGDATA=/home

upvoted 3 times

🗨️ 👤 **CAguys** 4 years, 3 months ago

First box is correct but the second box should be 'true'. Please refer the link which is given in the explanation!

upvoted 1 times

You are developing a web app that is protected by Azure Web Application Firewall (WAF). All traffic to the web app is routed through an Azure Application

Gateway instance that is used by multiple web apps. The web app address is contoso.azurewebsites.net.

All traffic must be secured with SSL. The Azure Application Gateway instance is used by multiple web apps.

You need to configure the Azure Application Gateway for the web app.

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

NOTE: Each correct selection is worth one point.

- A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.
- B. Convert the web app to run in an Azure App service environment (ASE).
- C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.
- D. In the Azure Application Gateway's HTTP setting, set the value of the Override backend path option to contoso22.azurewebsites.net.

**Suggested Answer: AD**

D: The ability to specify a host override is defined in the HTTP settings and can be applied to any back-end pool during rule creation.

The ability to derive the host name from the IP or FQDN of the back-end pool members. HTTP settings also provide an option to dynamically pick the host name from a back-end pool member's FQDN if configured with the option to derive host name from an individual back-end pool member.

A (not C): SSL termination and end to end SSL with multi-tenant services.

In case of end to end SSL, trusted Azure services such as Azure App service web apps do not require whitelisting the backends in the application gateway.


Therefore, there is no need to add any authentication certificates.


### Add HTTP setting


saiappgw-appgw

\* Protocol

☐ HTTP ☒ HTTPS


 Authentication certificates are not required for trusted Azure certificates for end to end ssl to work

\* Port 


443 

\* Request timeout (seconds)

20

Override backend path 

☒ Use for App service

☒ Use custom probe 

OK

Reference:

<https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview>

🗳️ 👤 **dirtygooback** Highly Voted 4 years, 7 months ago

A&D are correct. The question was not copied correctly from AZ-203

<https://www.examttopics.com/discussions/microsoft/view/21243-exam-az-203-topic-24-question-22-discussion/>

upvoted 52 times

🗳️ 👤 **iiiihhh** 4 years, 6 months ago

There is a difference between AZ-203 and this question. The AZ-203 question talks about contoso22. ... (and has consto22.... also in the pertinent answer), but this question has constoso. ... in the question and consto22 in the answer. This discrepancy makes the greatest confusion

upvoted 21 times

🗳️ 👤 **Bartimaeus** 3 years, 4 months ago

Nope, there isn't anything about App Service - you can't use it with normal web app.

C & D - you need both

upvoted 1 times

🗳️ 👤 **clarionprogrammer** 4 years, 2 months ago

Agreed A & D are correct.

Note: [T]here is no requirement to add authentication certificates when App service is selected as backend.

<https://docs.microsoft.com/en-us/azure/application-gateway/application-gateway-web-app-overview#tls-termination-and-end-to-end-tls-with-multi-tenant-services>

upvoted 7 times

🗳️ 👤 **samraw83** 3 years ago

Quite odd this should be AZ-104 question, It should be C & D, you have to update the secure certificate and have to update the backend pools as updated here:

<https://docs.microsoft.com/en-us/azure/app-service/configure-custom-container?pivots=container-linux>

upvoted 1 times

🗳️ 👤 **mrqc** Highly Voted 4 years, 5 months ago

Is this a valid question for AZ204? I don't see this topic in the skills measured?

upvoted 21 times

🗳️ 👤 **Krishnait06** Most Recent 5 months, 3 weeks ago

Selected Answer: AD

There is a option to (D) Override backend path and (A) Use for App Service while add HTTP Setting

upvoted 1 times

🗳️ 👤 **Faolba** 9 months, 1 week ago

Selected Answer: AC

A and C are correct.

upvoted 1 times

🗳️ 👤 **all\_cloud** 1 year, 2 months ago

question seems to be wrong

upvoted 2 times

🗳️ 👤 **Chriscrea8** 1 year, 3 months ago

From ChatGPT:

The correct actions to configure the Azure Application Gateway for the web app are:

A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting - This ensures that the Application Gateway can work seamlessly with the Azure Web App.

C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway - This is necessary to secure the traffic with SSL.

So, the correct answers are A and C.

upvoted 1 times



🗨️ 👤 **FeriAZ** 1 year, 4 months ago

Selected Answer: AC

C. For SSL traffic, the Azure Application Gateway needs to authenticate itself to the backend web app using an SSL certificate. Adding an authentication (or backend) certificate for contoso.azurewebsites.net to the Application Gateway is necessary to establish a secure SSL connection between the Application Gateway and the web app. This certificate ensures that the traffic remains encrypted and secure as it passes through the Application Gateway to the web app.

upvoted 2 times

🗨️ 👤 **manopeydaon** 1 year, 5 months ago

A and C,

The correct actions to configure Azure Application Gateway for the web app are:

- A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.
- C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.

Explanation:

A. Enabling the "Use for App service" setting in the HTTP setting of Azure Application Gateway ensures that the gateway can be used with Azure App Service, which includes the web app.

C. Adding an authentication certificate for contoso.azurewebsites.net is necessary to secure the traffic with SSL. This involves associating an SSL certificate with the custom domain (contoso.azurewebsites.net) used by the web app.

Therefore, options A and C are the correct actions.

upvoted 2 times

🗨️ 👤 **11\_NickName\_11** 1 year, 7 months ago

Not sure, but looks like there is no need for cert.

Application Gateway operates as an application delivery controller (ADC). It offers Transport Layer Security (TLS), previously known as Secure Sockets Layer (SSL), termination, cookie-based session affinity, round-robin load distribution, content-based routing, ability to host multiple websites, and security enhancements.

Application Gateway security enhancements include TLS policy management and end-to-end TLS support. Application security is strengthened by WAF integration into Application Gateway. The combination protects your web applications against common vulnerabilities. And it provides an easy-to-configure central location to manage.

<https://learn.microsoft.com/en-us/azure/web-application-firewall/ag/ag-overview#benefits>

upvoted 2 times

🗨️ 👤 **NicolasMous** 1 year, 9 months ago

Selected Answer: AC

A, C, Bings explanation: The correct answers are:

A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting. C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.

Explanation:

Option A: The Azure Application Gateway needs to be configured to use the App service setting. This is because the traffic is being routed through an Azure Application Gateway instance that is used by multiple web apps1.

Option C: Since all traffic must be secured with SSL, an authentication certificate for contoso.azurewebsites.net needs to be added to the Azure Application Gateway1.

Option B is not correct because converting the web app to run in an Azure App service environment (ASE) is not necessary for configuring the Azure Application Gateway for the web app1.

Option D is not correct because changing the backend path to a different web app (contoso22.azurewebsites.net) does not align with the requirement of routing traffic to contoso.azurewebsites.net1.

upvoted 5 times

🗨️ 👤 **Vladimir\_Gajinov** 1 year, 9 months ago

Selected Answer: AC

The right answers are A and C.

To configure the Azure Application Gateway for the web app hosted at contoso.azurewebsites.net and ensure that all traffic is secured with SSL, you should perform the following actions:

A. In the Azure Application Gateway's HTTP setting, enable the Use for App service setting.

This setting ensures that the Azure Application Gateway can communicate with the Azure App Service (web app) using the appropriate settings and protocols.

C. Add an authentication certificate for contoso.azurewebsites.net to the Azure Application Gateway.

This step is essential to enable SSL termination at the Azure Application Gateway. You need to upload and configure the SSL certificate for contoso.azurewebsites.net so that the gateway can securely handle HTTPS traffic for your web app.

The other options, B and D, are not necessary for securing traffic with SSL.

upvoted 2 times

  **MikeAWS** 1 year, 9 months ago

is this explanation from ChatGPT?

upvoted 1 times

  **tarek0811** 1 year, 11 months ago

it should be A and C, its in the doc:

The SSL certificate is needed to secure the traffic to the web app by encrypting the data between the client and the server.

If you do not add an SSL certificate to the Azure Application Gateway, then all traffic to the web app will be unencrypted. This means that anyone who can intercept the traffic can see the data that is being transmitted, including passwords, credit card numbers, and other sensitive information.

<https://learn.microsoft.com/en-us/azure/application-gateway/end-to-end-ssl-portal>

upvoted 1 times

  **[Removed]** 1 year, 11 months ago

D is definitely incorrect

upvoted 1 times

  **Mark16** 2 years ago

Selected Answer: AC

Should be AC:

The correct actions to configure the Azure Application Gateway for the web app are:

A. In the Azure Application Gateway's HTTP setting, enable the "Use for App service" setting. This ensures that the Application Gateway is configured to work properly with the web app running in Azure App Service.

C. Add an authentication certificate for "contoso.azurewebsites.net" to the Azure Application Gateway. This certificate is necessary to enable SSL/TLS encryption for securing the traffic between the clients and the Application Gateway.

Therefore, the correct options are A and C.

upvoted 3 times

  **Mark16** 2 years ago

Also for the people that answered D:



Option D, which suggests setting the value of the "Override backend path" option to "contoso22.azurewebsites.net," is not necessary or relevant in this scenario. The goal is to configure the Azure Application Gateway to secure the traffic for the web app hosted at "contoso.azurewebsites.net," not "contoso22.azurewebsites.net." Therefore, option D is not a correct action for this particular requirement.

upvoted 3 times

  **uffuchsi** 2 years, 4 months ago

Received this in my exam today (22/02/2023). Selected A and D. Score 927.

upvoted 4 times

  **jrnjerin** 2 years, 6 months ago

Nowhere it is mentiones contoso22

upvoted 3 times

  **hubekpeter** 2 years, 7 months ago

**Selected Answer: AD**

A. yes, app is using default appservice domain azurewebsites.net, so my assumption is, it's hosted as appservice, B - it's not needed as it's already running as App Service, therefore C - is not needed as well, D - yes you need either pool or redirection, in this case we need to redirect the traffic, there's 100% typo in hostname

upvoted 1 times

  **hubekpeter** 2 years, 7 months ago

Here's an article where the answer D is explained in detail. <https://learn.microsoft.com/en-us/azure/application-gateway/configure-web-app?tabs=customdomain%2Cazure-portal>

upvoted 3 times

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 develop a software as a service (SaaS) offering to manage photographs. Users upload photos to a web service which then stores the photos in Azure

Storage Blob storage. The storage account type is General-purpose V2.

When photos are uploaded, they must be processed to produce and save a mobile-friendly version of the image. The process to produce a mobile-friendly version of the image must start in less than one minute.

You need to design the process that starts the photo processing.

Solution: Use the Azure Blob Storage change feed to trigger photo processing.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

The change feed is a log of changes that are organized into hourly segments but appended to and updated every few minutes. These segments are created only when there are blob change events that occur in that hour.

Instead catch the triggered event, so move the photo processing to an Azure Function triggered from the blob upload.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed> <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

Community vote distribution

B (100%)

 **finnishr**  3 years, 10 months ago



The answer is 100% correct.

Change feed publishes records to the log within few minutes of the change. The process to produce a mobile-friendly version of the image must start in less than one minute.

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal>  
upvoted 23 times

 **AhmedAbdelAziz** 1 year, 4 months ago

Change feed is for logging purpose only  
upvoted 2 times

 **Zezere**  7 months ago

**Selected Answer: B**

So this question comes several times. The possibilities are:

1. Trigger the photo processing from Blob storage events.
2. Convert the Azure Storage account to a BlockBlobStorage storage account.
3. Move photo processing to an Azure Function triggered from the blob upload
4. Create an Azure Function app that uses the Consumption hosting model and that is triggered from the blob upload.
5. Use the Azure Blob Storage change feed to trigger photo processing

1. is not right because it does not involve using azure function -> it can take up to 10min
2. is not right. The type of the blob has no relevance here
3. is correct. Using Azure function implies low latency
4. Also correct. It is possible to use the consumption plan to reduce the costs.
5. Not correct. Change Feed is for the logs and has no relevance here.

upvoted 7 times

 **space\_cadet**  5 months, 4 weeks ago

**Selected Answer: B**

I haven't seen the right solution to this ask yet.

An event from the upload of a file to the web app should trigger a function app that processes the images and has an output binding to upload to a blob container.

So an event grid topic might be the solution

upvoted 1 times

🗲️ 👤 **AhmedAbdelAziz** 1 year, 4 months ago

Change feed is for logging and took more than min , we should use EventGrid with Azure functions

upvoted 1 times

🗲️ 👤 **FeriAZ** 1 year, 4 months ago

**Selected Answer: B**

Azure Functions with a Blob Storage trigger is better suited for this requirement.

upvoted 1 times

🗲️ 👤 **1989jay** 1 year, 5 months ago

This question was on exam. 09/01/2024. Passed 872.

upvoted 3 times

🗲️ 👤 **Ciupaz** 1 year, 7 months ago

**Selected Answer: B**

No. Instead use Azure Event Grid.

upvoted 1 times

🗲️ 👤 **Vladimir\_Gajinov** 1 year, 10 months ago

**Selected Answer: B**

Simply NO!

upvoted 1 times

🗲️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 3 times

🗲️ 👤 **chirandev** 2 years, 3 months ago

**Selected Answer: B**

Change feed is for logging.

upvoted 3 times

🗲️ 👤 **Esward** 2 years, 5 months ago

Given answer is correct!

upvoted 1 times

🗲️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

correct.

upvoted 1 times

🗲️ 👤 **gmishra88** 2 years, 8 months ago

There is blob storage events through event grid and the azure-function-triggers as alternatives for faster processing

upvoted 1 times

🗲️ 👤 **gmishra88** 2 years, 8 months ago

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed?tabs=azure-portal#should-i-use-the-change-feed-or-storage-events>

upvoted 1 times

🗲️ 👤 **Ganesh1983** 3 years, 1 month ago

**Selected Answer: B**

change feed is more to do with audit requirements.

upvoted 1 times

🗲️ 👤 **Freidrich** 3 years, 4 months ago

**Selected Answer: B**

The correct answer is B: No.



upvoted 1 times

  **edengoforit** 3 years, 5 months ago

same question here and the answer is No according to the question below

<https://www.examttopics.com/discussions/microsoft/view/26699-exam-az-204-topic-1-question-19-discussion/>

upvoted 2 times

  **r\_k** 3 years, 7 months ago

Important info related to this question. <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-trigger?tabs=csharp>

upvoted 1 times

You are developing a web application that runs as an Azure Web App. The web application stores data in Azure SQL Database and stores files in an Azure Storage account. The web application makes HTTP requests to external services as part of normal operations. The web application is instrumented with Application Insights. The external services are OpenTelemetry compliant. You need to ensure that the customer ID of the signed in user is associated with all operations throughout the overall system. What should you do?


- A. Add the customer ID for the signed in user to the CorrelationContext in the web application
- B. On the current SpanContext, set the TraceId to the customer ID for the signed in user
- C. Set the header Ocp-Apim-Trace to the customer ID for the signed in user
- D. Create a new SpanContext with the TraceFlags value set to the customer ID for the signed in user

**Suggested Answer: A**

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/correlation>

Community vote distribution



**1989jay** Highly Voted 11 months, 3 weeks ago

This question was on exam. 09/01/2024. Passed 872.  
upvoted 13 times

**neelkanths** Highly Voted 8 months ago

**Selected Answer: A**

Got it on 20 April 2024...Marks > 950...answer is correct....all questions from examtopics 400 question bank  
upvoted 9 times

**raja33** Most Recent 2 months, 1 week ago

**Selected Answer: A**

Got this in the exam today! Apr 19, 2025  
upvoted 1 times

**applepie** 1 year, 5 months ago

got this question today, answer A - 7/30/2023, score 895/1000  
upvoted 5 times

**imanonion** 1 year, 11 months ago

**Selected Answer: A**

Answer by ChatGPT, which I think looks legit:

"I would choose option A, Add the customer ID for the signed in user to the CorrelationContext in the web application.

The CorrelationContext is a way to associate contextual information with a request as it flows through the system. It allows you to track a request as it passes through different components of the system, and to identify related log entries and telemetry data. By adding the customer ID to the CorrelationContext in the web application, you can ensure that it is associated with all operations throughout the overall system. This will allow you to track the request and identify related log entries and telemetry data for a specific customer.

upvoted 4 times

**imanonion** 1 year, 11 months ago

Option B, On the current SpanContext, set the TraceId to the customer ID for the signed in user, is not correct because the TraceId is a unique identifier for a trace, and it is not meant to be used to store contextual information about the request.

Option C, Set the header Ocp-Apim-Trace to the customer ID for the signed in user, is not correct because the Ocp-Apim-Trace header is used by Azure API Management to track requests as they pass through the API gateway, and it is not related to application instrumentation or tracing.

Option D, Create a new SpanContext with the TraceFlags value set to the customer ID for the signed in user, is not correct because the TraceFlags value is used to indicate the state of the trace and is not meant to be used to store contextual information about the request."

upvoted 7 times

  **Puja079888** 10 months, 1 week ago

good explanation



upvoted 1 times

  **OPT\_001122** 2 years, 1 month ago

**Selected Answer: A**

A. Add the customer ID for the signed in user to the CorrelationContext in the web application

upvoted 2 times

  **Akiu** 2 years, 2 months ago

OpenTelemetry uses correlation context and it fits the requirements:

Correlation context carries user-defined properties. These properties are typically data that you would like to eventually aggregate for correlation analysis or use to filter your trace data, such as a customer identifier, ...

<https://lightstep.com/opentelemetry/context-propagation>

upvoted 3 times

  **chessace2000** 2 years, 2 months ago

**Selected Answer: B**

It has to be OpenTelemetry compliant. This means using W3C correlation spec. CorrelationContext (A) is legacy correlation spec not W3C.

upvoted 2 times

  **warchoon** 1 year, 10 months ago

tracestate: Carries system-specific tracing context.

It is a part of correlation headers using W3C TraceContext.

Correlation-Context is deprecated. But this is the same as the Request-Context which is currently in use.

upvoted 1 times

  **willchenxa** 2 years, 3 months ago

**Selected Answer: A**

A SpanContext represents the portion of a Span which must be serialized and propagated along side of a Baggage.



<https://open-telemetry.github.io/opentelemetry-js-api/interfaces/spancontext.html>

upvoted 5 times

  **warchoon** 1 year, 10 months ago

Agree <https://learn.microsoft.com/en-us/azure/azure-monitor/app/opentelemetry-enable?tabs=java#set-the-user-id-or-authenticated-user-id>

upvoted 1 times

  **finnishr** 2 years, 4 months ago

**Selected Answer: B**

The correct answer is B.

upvoted 1 times

  **DonH** 1 year, 6 months ago


This answer is useless without any explanation. Come on guys this is annoying and NOT helpful. Please clarify WHY it is the correct answer. Now we have a group of people stating A is correct and an other group stating B is correct. How are we suppose to know what is correct?

upvoted 3 times

  **lhy904** 2 years, 7 months ago

Why not B?

upvoted 2 times

  **warchoon** 1 year, 10 months ago

"The value of trace-id SHOULD be globally unique. One recommended method to ensure global uniqueness, as well as to address some privacy and security considerations, to a satisfactory degree of certainty is to randomly (or pseudo-randomly) generate the trace-id. Implementers SHOULD use a trace-id generation method which randomly (or pseudo-randomly) generates at least the right-most 7 bytes of the ID." by specs ChatGPT is lying (as usual) that you can't use it to request, but you can't use Traceld with all operations throughout the overall system

upvoted 2 times

  **nt\_ar** 2 years, 8 months ago

**Selected Answer: A**

That's correct

upvoted 2 times



**HOTSPOT -**

You are developing an Azure Function App. You develop code by using a language that is not supported by the Azure Function App host. The code language supports HTTP primitives.

You must deploy the code to a production Azure Function App environment.

You need to configure the app for deployment.

Which configuration values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area****Configuration parameter****Configuration value**

Publish

	▼
Code	
Docker Container	

Runtime stack

	▼
Node.js	
Python	
PowerShell Core	
Custom Handler	

Version

	▼
14 LTS	
7.0	
custom	

**Answer Area****Configuration parameter****Configuration value**

Publish

	▼
Code	
Docker Container	

Runtime stack

	▼
Node.js	
Python	
PowerShell Core	
Custom Handler	

Version

	▼
14 LTS	
7.0	
custom	

Suggested Answer:

A custom handler can be deployed to every Azure Functions hosting option. If your handler requires operating system or platform dependencies (such as a language runtime), you may need to use a custom container. You can create and deploy your code to Azure Functions as a custom Docker container.

Box 2: PowerShell core -

When creating a function app in Azure for custom handlers, we recommend you select .NET Core as the stack. A "Custom" stack for custom handlers will be added in the future.

PowerShell Core (PSC) is based on the new .NET Core runtime.

Box 3: 7.0 -

On Windows: The Azure Az PowerShell module is also supported for use with PowerShell 5.1 on Windows.

On Linux: PowerShell 7.0.6 LTS, PowerShell 7.1.3, or higher is the recommended version of PowerShell for use with the Azure Az PowerShell module on all platforms.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-create-function-linux-custom-image> <https://docs.microsoft.com/en-us/powershell/azure/install-az-ps?view=azps-7.1.0>

🗳️ 👤 **SaintBahamut** Highly Voted 3 years, 2 months ago

answer is wrong: if you pick docker container you cant specify stack and version,

So ill go with

>Code

>Custom Handler

>custom (only option when you pick Custom Handler)

upvoted 114 times

🗳️ 👤 **Esward** 2 years, 5 months ago

Agreed

upvoted 1 times

🗳️ 👤 **JustHereToLearn** 2 years, 6 months ago

To further support this answer see this video tutorial from Microsoft official channel where they also select "Custom Handler" and do not use Docker: <https://www.youtube.com/watch?v=roWbjuo-c30>

upvoted 3 times

🗳️ 👤 **mihailos** 1 year, 10 months ago

Got it in exam on 28/08/23. Went with this answer. Scored 912

upvoted 4 times

🗳️ 👤 **hubekpeter** 2 years, 7 months ago

Sorry first answer should be docker

upvoted 3 times

🗳️ 👤 **EricCao** Highly Voted 2 years, 8 months ago

I checked it on Azure portal and the answer should be

Code > Custom Handler > custom

upvoted 19 times

🗳️ 👤 **JamieS** 2 years, 5 months ago

Just tried it to, he's right., Code > Custom Handler > Custom

upvoted 2 times

🗳️ 👤 **warchoon** 2 years, 4 months ago

Did you use a language that is NOT supported by the Azure Function App host?

upvoted 2 times

🗳️ 👤 **Thameur01** Most Recent 2 months, 1 week ago

please fix the answer, it just confusing:

- Code

- Custom Handler

- Custom

upvoted 1 times

🗳️ 👤 **AryaNiaz** 10 months ago

When you select Code as the publish option, Azure expects the code to be in a language that the Azure Functions runtime natively supports (such as C#, JavaScript, Python, Java, PowerShell, etc.). Since your language is not supported natively, you would not be able to use the Code option to deploy your function successfully.

upvoted 1 times

🗨️ 👤 **royalbaby** 11 months, 4 weeks ago

Passed the exam. Got this Q on 07/07.

Case Study - City Power & Lights

Code > Custom Handler - Custom

upvoted 3 times

🗨️ 👤 **giuliohome** 12 months ago

if you use a docker container then the runtime and all is defined in the Dockerfile, hence I think it is not docker here because you have a "custom handler" and likely "code" makes more sense as an answer.

upvoted 1 times

🗨️ 👤 **renewss** 1 year ago

Custom handlers are lightweight web servers that receive events from the Azure Functions host. Any language that supports HTTP primitives can implement a custom handler. This means that custom handlers can be used to create functions in languages that aren't officially supported. To learn more, see Azure Functions custom handlers.

<https://www.examtopycs.com/exams/microsoft/az-204/view/8/>

upvoted 1 times

🗨️ 👤 **onlyforheros** 1 year, 3 months ago

Got it in exam on 13.03.2024. Score: 910.

upvoted 1 times

🗨️ 👤 **JeganHalifax** 1 year, 3 months ago

I got this question - 09/03/224

upvoted 1 times

🗨️ 👤 **arunkum** 1 year, 6 months ago

Got it in the exam 14/12/23. Went with given answer. All questions are from ExamTopics. Case study - VanArsdel, Ltd (11 questions)

upvoted 1 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 6 months ago

I got this question on my exam, 2023Dec, go with what I remember was the most voted answer. Score 902, most of the questions were here, slightly different on wording because the Azure Ad <-> Entra Id change. Case was City Power & Light. Good luck!

upvoted 1 times

🗨️ 👤 **Veeresh114** 1 year, 5 months ago

Does exam topics include case studies in these 200 questions or need to buy premium?

upvoted 3 times

🗨️ 👤 **NPE\_** 1 year, 7 months ago

Got this on the exam Nov. 07 2023

upvoted 2 times

🗨️ 👤 **11\_NickName\_11** 1 year, 7 months ago

Have you passed?

Which one you choose: Code > Custom Handler > Custom?

upvoted 1 times

🗨️ 👤 **richardc7** 1 year, 9 months ago

got in the exam 28/09/2023

upvoted 1 times

🗨️ 👤 **RuffBoii** 1 year, 9 months ago

Had this on my exam today.

upvoted 1 times

🗨️ 👤 **nikipediaa** 1 year, 9 months ago

me too, today

upvoted 1 times

🗨️ 👤 **Tarajee** 1 year, 9 months ago

Got this on exam 2023sept  
upvoted 2 times

🗨️ 👤 **nlk0** 1 year, 10 months ago

Got this today 26/08/2023, went with the most voted. 880  
Case study was contoso -> 9 questions  
+ 35 question, two series.  
3 of them were not from examtopics, all others were  
upvoted 4 times

🗨️ 👤 **LittleBean** 1 year, 6 months ago

did you bought contributor access?  
upvoted 3 times

🗨️ 👤 **kvtborad** 1 year, 10 months ago

I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.  
Code, Custom Handler, Custom  
upvoted 3 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

Thank you for informing us! Glad to hear questions are from here.  
upvoted 2 times

## DRAG DROP -

You provision virtual machines (VMs) as development environments.

One VM does not start. The VM is stuck in a Windows update process. You attach the OS disk for the affected VM to a recovery VM.

You need to correct the issue.

In which order should you perform the actions? 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

Run the following command at an elevated command prompt:

```
dism /image:\ /get-packages > c:\temp\Patch.txt
```

Run the following command at an elevated command prompt:

```
dism /Image:<Attached OS disks>:\ /Remove  
Package /PackageName:<package name to delete>
```

Detach the OS disk and recreate the VM

Open C:\temp\Patch.txt file and locate the update that is in a pending state

## Answer Area



## Suggested Answer:

## Actions

## Answer Area

Run the following command at an elevated command prompt:

```
dism /image:\ /get-packages > c:\temp\Patch.txt
```

Open C:\temp\Patch.txt file and locate the update that is in a pending state

Run the following command at an elevated command prompt:

```
dism /Image:<Attached OS disks>:\ /Remove  
Package /PackageName:<package name to delete>
```

Detach the OS disk and recreate the VM



Remove the update that causes the problem

1. Take a snapshot of the OS disk of the affected VM as a backup.
2. Attach the OS disk to a recovery VM.
3. Once the OS disk is attached on the recovery VM, run diskmgmt.msc to open Disk Management, and ensure the attached disk is ONLINE.
4. (Step 1) Open an elevated command prompt instance (Run as administrator). Run the following command to get the list of the update packages that are on the attached OS disk: `dism /image:<Attached OS disk>:\ /get-packages > c:\temp\Patch_level`
5. (Step 2) Open the C:\temp\Patch\_level.txt file, and then read it from the bottom up. Locate the update that's in Install Pending or Uninstall Pending state.
6. Remove the update that caused the problem:  
`dism /Image:<Attached OS disk>:\ /Remove-Package /PackageName:<PACK`
7. (Step 4) Detach the OS disk and recreate the VM. Then check whether the issue is resolved.

Reference:

<https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-stuck-updating-boot-error>

nedlo Highly Voted 2 years, 3 months ago

sure i will use this knowledge for my developer work...

upvoted 48 times

🗨️ 👤 **grudzin** 1 year, 12 months ago

Question says "The VM is stuck in a Windows update process". I ques it's common ;)  
upvoted 8 times

🗨️ 👤 **br00net** 2 years, 1 month ago

clearly, every developer needs to memorize such knowledge instead of searching in the documentation  
upvoted 7 times

🗨️ 👤 **85e4c91** 1 year, 8 months ago

I guess developers don't have access to google nowadays  
upvoted 5 times

🗨️ 👤 **Magis** Highly Voted 🍌 2 years, 7 months ago

Correct.

All steps are listed here - <https://docs.microsoft.com/en-us/troubleshoot/azure/virtual-machines/troubleshoot-stuck-updating-boot-error>

upvoted 16 times

🗨️ 👤 **Christian\_garcia\_martin** Most Recent 🔍 11 months, 1 week ago

Is this a developer responsibility?  
upvoted 1 times

🗨️ 👤 **Lucifer14** 2 years, 6 months ago

In 28-12-2022 exam  
upvoted 3 times

🗨️ 👤 **memadk** 2 years, 8 months ago

Is this a part of AZ-204? Then I must have missed something....  
upvoted 5 times

🗨️ 👤 **martijn44** 2 years, 8 months ago

I doubt it.. it seems many of these questions were copied over from AZ-203.  
upvoted 3 times

🗨️ 👤 **finnishr** 2 years, 10 months ago

100% correct!!  
upvoted 4 times

🗨️ 👤 **SoftSol** 3 years, 2 months ago

Correct  
upvoted 2 times

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 develop an HTTP triggered Azure Function app to process Azure Storage blob data. The app is triggered using an output binding on the blob. The app continues to time out after four minutes. The app must process the blob data.

You need to ensure the app does not time out and processes the blob data.

Solution: Update the functionTimeout property of the host.json project file to 10 minutes.

Does the solution meet the goal?

- A. Yes
- B. No

**Suggested Answer: B**

Instead pass the HTTP trigger payload into an Azure Service Bus queue to be processed by a queue trigger function and return an immediate HTTP success response.

Note: Large, long-running functions can cause unexpected timeout issues. General best practices include:

Whenever possible, refactor large functions into smaller function sets that work together and return responses fast. For example, a webhook or HTTP trigger function might require an acknowledgment response within a certain time limit; it's common for webhooks to require an immediate response. You can pass the HTTP trigger payload into a queue to be processed by a queue trigger function. This approach lets you defer the actual work and return an immediate response.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

Community vote distribution

B (100%)


 **stfnmrr** Highly Voted 2 years, 9 months ago

Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>  
upvoted 8 times

 **Ganesh1983** Highly Voted 3 years, 1 month ago


**Selected Answer: B**


HTTP triggers on function apps can timeout in 230 seconds. Hence B is the answer.  
upvoted 5 times

 **surprise0011** 2 years, 2 months ago  
received 2023-04-17 went given answer, score 926  
upvoted 3 times

 **Nitin23** Most Recent 1 year, 9 months ago

Answer is A as per chatgpt.  
Its a valid solution  
upvoted 1 times

 **MDWPartners** 1 year, 9 months ago  
Do not use chatGPT as a source, is really faulty with this kind of questions  
upvoted 10 times

 **dddddd111** 1 year, 7 months ago  
I did ask chatgpt and said it is "yes", but I asked to check the link provided for function scale timeout. Then chatgpt changed answer to "no".

No, updating the functionTimeout property in the host.json project file to 10 minutes won't meet the goal. Even if you increase the function timeout, the Azure Load Balancer has an idle timeout of 230 seconds for HTTP-triggered functions. As a result, the function will still time out after 230 seconds.  
upvoted 3 times

🗨️ 👤 **Tej\_007** 9 months, 4 weeks ago

absolutely correct.

upvoted 1 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

Received this in my exam today (22/02/2023). Selected B. Score 927.

upvoted 3 times

🗨️ 👤 **Esward** 2 years, 5 months ago

B is correct answer

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

upvoted 1 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-best-practices>

upvoted 1 times

🗨️ 👤 **sghaha** 3 years, 2 months ago

Why not a?

upvoted 1 times



## HOTSPOT -

You are developing an Azure Durable Function based application that processes a list of input values. The application is monitored using a console application that retrieves JSON data from an Azure Function diagnostic endpoint.

During processing a single instance of invalid input does not cause the function to fail. Invalid input must be available to the monitoring application.

You need to implement the Azure Durable Function and the monitoring console application.

How should you complete the code segments? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

```
[FunctionName("App")]
public static async Task<List<string>> RunOrchestrator(
    [OrchestrationTrigger] IDurableOrchestrationContext context) {
    EntityId[] input = . . .
    int errIndex = . . .

    ;
    context.SetOutput(input[errIndex])
    context.SetCustomStatus(input[errIndex])
    context.SignalEntity(input[errIndex], "error")
    await context.CallEntityAsync(input[errIndex], "error")
}

using (var client = new HttpClient())
{
    while (true)
    {
        var response = await client.GetAsync(". . .");
        response.EnsureSuccessStatusCode();
        var json = await response.Content.ReadAsStringAsync();
        dynamic result = JsonConvert.DeserializeObject(json);
        if (result.runtimeStatus == "
        Failed
        Awaited
        Listening
        Completed
    )
    {
        return result.
        input
        output
        runtimeStatus
        customStatus
    }
}
```

## Answer Area

```
[FunctionName("App")]
public static async Task<List<string>> RunOrchestrator(
    [OrchestrationTrigger] IDurableOrchestrationContext context) {
    EntityId[] input = . . .
    int errIndex = . . .

    ;
    context.SetOutput(input[errIndex])
    context.SetCustomStatus(input[errIndex])
    context.SignalEntity(input[errIndex], "error")
    await context.CallEntityAsync(input[errIndex], "error")
}

using (var client = new HttpClient())
{
    while (true)
    {
        var response = await client.GetAsync(". . .");
        response.EnsureSuccessStatusCode();
        var json = await response.Content.ReadAsStringAsync();
        dynamic result = JsonConvert.DeserializeObject(json);
        if (result.runtimeStatus == "
        Failed
        Awaited
        Listening
        Completed
        ")
        {
            return result.
            input
            output
            runtimeStatus
            customStatus
        }
    }
}
```

Suggested Answer:

Box 1: await context.CallEntityAsync(input[errindex],"error")

Orchestration signals and calls an entity

Orchestrator functions can access entities by using APIs on the orchestration trigger binding.

Example:

```
[FunctionName("CounterOrchestration")]
public static async Task Run(
    [OrchestrationTrigger] IDurableOrchestrationContext context)
{
    var entityId = new EntityId(nameof(Counter), "myCounter");
    // Two-way call to the entity which returns a value - awaits the response
    int currentValue = await context.CallEntityAsync<int>(entityId, "Get");
}
```

Box 2: Failed -

During processing a single instance of invalid input does not cause the function to fail.

Note: RuntimeStatus: One of the following values:

Failed: The instance failed with an error.

Completed: The instance has completed normally.

Terminated: The instance was stopped abruptly.

Pending: The instance has been scheduled but has not yet started running.

Running: The instance has started running.

ContinuedAsNew: The instance has restarted itself with a new history. This state is a transient state.

Box 3: Input -

Invalid input must be available to the monitoring application.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-entities> <https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-instance-management>

🗨️ 👤 **RochaG2** Highly Voted 2 years, 10 months ago

I've not seen any of this stuff in Microsoft's Learning Path's for AZ-204

upvoted 41 times

🗨️ 👤 **OPT\_001122** 2 years, 9 months ago

yes. there are a few questions added in last month. not sure if all of them are of AZ-204

upvoted 2 times

🗨️ 👤 **gmishra88** 2 years, 8 months ago

The learning path and the preparation videos are showing things you need not know for the exam.

upvoted 22 times

🗨️ 👤 **130nk3r5** 1 year, 6 months ago

Got this today.

Went with answer here.

Score 927

upvoted 6 times

🗨️ 👤 **Akiu** Highly Voted 2 years, 8 months ago

I'd say:

SetOutput()

== Completed

result.output

The function needs to return the one invalid input, so you could use either SetOutput or SetCustomStatus, but custom status is not available after the function completes.

The one invalid input does not cause the function to fail, therefore Completed.

upvoted 27 times

🗨️ 👤 **Christian\_garcia\_martin** 10 months, 2 weeks ago

last one should be runtimeStatus

upvoted 1 times

🗨️ 👤 **P000JAAAAAAAAA** 2 years, 6 months ago

is this correct ?

upvoted 4 times

🗨️ 👤 **regex** 1 year, 8 months ago

async functions needs an await part. SetOutput() won't qualify for that as function.

upvoted 1 times

🗨️ 👤 **gebpt** 1 year ago

you've got await in "var json = await response.Content.ReadAsStringAsync();" ,

upvoted 1 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

This one seems right!

upvoted 1 times

🗨️ 👤 **8ac3742** Most Recent 11 months ago

I choose setStatus/Completed/customStatus, setStatus make invalid input available to monitoring,

response.EnsureSuccessStatusCode() would throw Exception if the result.runtimeStatus != "Complete" so I choose result.runtimeStatus == "Completed" and return result.customStatus.

upvoted 1 times

🗨️ 👤 **oskx2** 1 year, 3 months ago

The output of the function has to match the method return type. It cannot be SetOutput. You can use CustomStatus, since the other 2 do not make sense. And the output is available on completion since the function does not throw an exception. But because the function return type is List<string>, and it may be used for another purpose, we use the custom status to get the invalid input.

Also, the custom status is available after the function completes.

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-api#get-instance-status>

upvoted 2 times

🗳️ 👤 **CarlosTheBoldest** 1 year, 6 months ago

I check it and `setCustomStatus` makes sense, cause they want to show the input that is causing errors (see link below). Apart of that, the `ensuresuccessstatuscode` will throw an exception if it fails, so I would go with "completed" and output

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-custom-orchestration-status?tabs=csharp>

<https://learn.microsoft.com/en-us/dotnet/api/system.net.http.httpresponsemessage.ensuresuccessstatuscode?view=net-8.0>

upvoted 4 times

🗳️ 👤 **Vegetta95** 1 year, 6 months ago

1. `SetCustomStatus`. There is no point in calling `context.CallEntityAsync` with "error" as parameter before knowing if error occurred or not
2. Failed (only negative status)
3. `runTimeStatus` ( This is the only property we are sure that dynamic object possess. We don't know the custom response from the url)

upvoted 1 times

🗳️ 👤 **Ciupaz** 1 year, 6 months ago

Hope don't find this question in my exam.

upvoted 9 times

🗳️ 👤 **Vukile** 1 year, 7 months ago

On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors

upvoted 3 times

🗳️ 👤 **longnguyendh** 1 year, 11 months ago

I think that the answer is

`SetCustomStatus`

Failed

`result.customStatus`

upvoted 3 times

🗳️ 👤 **MysticalSam** 2 years ago

This question was in today's exam at 10-June-2023

upvoted 7 times

🗳️ 👤 **NavinD** 1 year, 11 months ago

How much percent questions did you get from here? Do you have contributor access?

upvoted 2 times

🗳️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 4 times

🗳️ 👤 **PP2015** 2 years, 3 months ago

Did my exam on 3/29/2023. This question was on it. Selected `context.Signal(input[errIndex], "error")`, Completed, output. Score 850

upvoted 2 times

🗳️ 👤 **red0400** 2 years, 3 months ago

Answer - `setOutput`, completed, output

upvoted 5 times

🗳️ 👤 **pkolodziej** 2 years, 3 months ago

For those who are wondering if it's a `SignalEntity`, it's not:

`SignalEntity(EntityId, DateTime, String, Object)`

Signals an operation to be performed by an entity at a specified time. Any result or exception is ignored (fire and forget).

We're clearly waiting for completion here, so have to use `CallEntity`,



upvoted 1 times

🗳️ 👤 **pkolodziej** 2 years, 3 months ago

I correct my mistake: it should be CustomStatus


<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-custom-orchestration-status?tabs=csharp#querying-custom-status-with-http>

upvoted 4 times

  **uffuchsi** 2 years, 4 months ago

Received this in my exam today (22/02/2023). Selected context.Signal(input[errIndex], "error"), Completed, output. Score 927.

upvoted 10 times

  **memobed890** 2 years, 3 months ago

De donde saco la información para decir que son esas respuestas?

upvoted 4 times

  **RealRaymond** 2 years, 4 months ago

Answer should be CustomStatus.

Ref: <https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-custom-orchestration-status?tabs=csharp#querying-custom-status-with-http>

upvoted 2 times

  **manopeydakon** 2 years, 4 months ago

Nigros the answer is correct since we are after the failed input values.

upvoted 5 times

You are developing an Azure Durable Function to manage an online ordering process.

The process must call an external API to gather product discount information.

You need to implement the Azure Durable Function.

Which Azure Durable Function types should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Orchestrator
- B. Entity
- C. Client
- D. Activity

#### Suggested Answer: AB


The Durable Functions extension exposes a set of built-in HTTP APIs that can be used to perform management tasks on orchestrations, entities, and task hubs.

These HTTP APIs are extensibility webhooks that are authorized by the Azure Functions host but handled directly by the Durable Functions extension.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-http-api>

Community vote distribution



**serpevi** 2 years, 9 months ago

Got this in 09/22 , went with Orchestrator and Entity, score 927.

upvoted 22 times

**ProtossOR89144** 12 months ago

" Like orchestrator functions, entity functions are functions with a special trigger type, entity trigger." - u can not call Entity from Orchestrator...  
right answer is Orchestrator and Activity

upvoted 1 times

**OPT\_001122** 2 years, 9 months ago

Thanks for mentioning the exam date in which it was asked.

upvoted 4 times

**[Removed]** 2 years, 5 months ago

Entity is for storing state which isn't called out as requirement. An activity can return/work with the HTTP req. Like mentioned before this, Dur  
Func 2.0 supports DurableHttpResponse. (Grats/Great score)

upvoted 6 times

**Chief2110** 2 years, 9 months ago

It should be orchestrator, client & activity right?

upvoted 20 times

**coffecold** 2 years, 8 months ago

yes, see <https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview>

upvoted 3 times

**Tej\_007** 9 months, 4 weeks ago

this link says orchestrator function can only be invoked by a client function.

even for testing purpose also it says to create a test client for testing the test orchestrator.

so if we are selecting orchestrator, client is also must right?

upvoted 1 times

**Esward** 2 years, 5 months ago

orchestrator and activity

upvoted 7 times

🗨️ 👤 **JustHereToLearn** 2 years, 5 months ago

Yes I think orchestrator and activity is correct.

Activity for this reason "Unlike orchestrator functions, activity functions aren't restricted in the type of work you can do in them. Activity functions are frequently used to make network calls or run CPU intensive operations. An activity function can also return data back to the orchestrator function."

Reference: <https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions>  
upvoted 3 times

🗨️ 👤 **TonyMel** 2 years, 3 months ago

correct, in 2023Mar24, score: 904/1000

upvoted 5 times

🗨️ 👤 **Tej\_007** Most Recent 9 months, 4 weeks ago

why can't the client function be part of the answer why only orchestrator and activity.

some one posted the link below <https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview>.

it clearly says an orchestrator function can only be invoked by a client function.

even for testing the orchestrator it says to create a test client function.

client is a must if we need to have an orchestrator. isn't it ?

upvoted 1 times

🗨️ 👤 **At09** 1 year ago

orchestrator and activity

upvoted 1 times

🗨️ 👤 **FeriAZ** 1 year, 4 months ago

Selected Answer: AD

Activity functions are the building blocks of the Durable Function that actually perform the work. They are called by the Orchestrator function. In this scenario, an Activity function would be responsible for calling the external API to gather product discount information. Because Activity functions can perform I/O operations, compute-heavy processing, or call external services, they are well-suited for this task.

upvoted 5 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

A and D,

To implement an Azure Durable Function for managing an online ordering process that calls an external API to gather product discount information, you should use the following Azure Durable Function types:

A. Orchestrator

D. Activity

Explanation:

Orchestrator functions define the control flow and are responsible for calling other functions (activities).

Activity functions perform a unit of work, such as calling an external API to gather product discount information.

So, the combination of Orchestrator (A) and Activity (D) functions is suitable for this scenario.

upvoted 2 times

🗨️ 👤 **Kanasan** 1 year, 10 months ago

Selected Answer: AD

A and D:

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions>

Activity functions are frequently used to make network calls or run CPU intensive operations. An activity function can also return data back to the orchestrator function.

upvoted 5 times

🗨️ 👤 **Kanasan** 1 year, 10 months ago

A and D:

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#activity-functions>

Activity functions are frequently used to make network calls or run CPU intensive operations. An activity function can also return data back to the orchestrator function.

upvoted 1 times

🗨️ 👤 **longnguyendh** 1 year, 11 months ago

It should be orchestrator & activity

Activity Function: This type of function is responsible for performing individual tasks, such as calling an external API to gather product discount information

upvoted 1 times

🗨️ 👤 **RayD7** 2 years, 1 month ago

Durable functions are out of scope as of 28th April 2023

upvoted 6 times

🗨️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

My case:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 1 times

🗨️ 👤 **Raja419** 2 years, 1 month ago

**Selected Answer: AD**

The "Orchestrator" function type is responsible for coordinating the workflow and managing the execution of the "Activity" function type. It receives inputs, calls external APIs, and makes decisions based on the results of the API calls.

The "Activity" function type is responsible for performing specific actions or tasks, such as calling an external API to gather product discount information. It receives inputs from the "Orchestrator" function and returns results back to it.

upvoted 2 times

🗨️ 👤 **Saluk\_DE** 2 years, 3 months ago

Question was in Exam 2023-03-30

upvoted 2 times

🗨️ 👤 **maqsoodshah** 2 years, 3 months ago

Correct Answers:

A. Orchestrator

D. Activity

Activity functions

Activity functions are the basic unit of work in a durable function orchestration. Activity functions are the functions and tasks that are orchestrated in the process. For example, you might create an orchestrator function to process an order. The tasks involve checking the inventory, charging the customer, and creating a shipment. Each task would be a separate activity function. These activity functions may be executed serially, in parallel, or some combination of both.

Entity functions

Entity functions define operations for reading and updating small pieces of state. We often refer to these stateful entities as durable entities. Like orchestrator functions, entity functions are functions with a special trigger type, entity trigger. They can also be invoked from client functions or from orchestrator functions

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview>

upvoted 3 times

🗨️ 👤 **adilkhan** 2 years, 3 months ago

A,D chatGPT

upvoted 1 times

🗨️ 👤 **DonH** 2 years ago

Again stop with this ChatGPT nonsense respons. ChatGPT is NOT a source rather than a search engine. It can also give incorrect answers. So stating its from ChatGPT doesn't say a thing about it's validity.


upvoted 8 times

🗨️ 👤 **CafeLago** 10 months, 3 weeks ago

ChatGPT is not a search engine. It can give incorrect answers, but so can you. Nobody trusts answers 100% here. You can ignore ChatGPT answers! Provide your correct answer instead of this rant.



upvoted 1 times

  **70PineApple** 2 years, 4 months ago

Got this in exam today..20/02/23

score: 817

upvoted 2 times

  **Priya0703** 2 years, 4 months ago

On exam today 20-02-2023

upvoted 2 times

## DRAG DROP -

You are authoring a set of nested Azure Resource Manager templates to deploy multiple Azure resources.

The templates must be tested before deployment and must follow recommended practices.

You need to validate and test the templates before deployment.

Which tools should you use? To answer, drag the appropriate tools to the correct requirements. Each tool may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Tools	Requirement	Tool
Parameter file	Determine whether the templates follow recommended practices.	Tool
Template function		
Azure Resource Manager test toolkit	Test and validate changes that templates will make to the environment.	Tool
User-defined function		
What-if operation		
Azure Deployment Manager		

## Suggested Answer:

Tools	Requirement	Tool
Parameter file	Determine whether the templates follow recommended practices.	Azure Resource Manager test toolkit
Template function		
Azure Resource Manager test toolkit	Test and validate changes that templates will make to the environment.	What-if operation
User-defined function		
What-if operation		
Azure Deployment Manager		

Box 1: Azure Resource Manager test toolkit

Use ARM template test toolkit -

The Azure Resource Manager template (ARM template) test toolkit checks whether your template uses recommended practices. When your template isn't compliant with recommended practices, it returns a list of warnings with the suggested changes. By using the test toolkit, you can learn how to avoid common problems in template development.

Box 2: What-if operation -

ARM template deployment what-if operation

Before deploying an Azure Resource Manager template (ARM template), you can preview the changes that will happen. Azure Resource Manager provides the what-if operation to let you see how resources will change if you deploy the template. The what-if operation doesn't make any changes to existing resources.

Instead, it predicts the changes if the specified template is deployed.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit> <https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if>

 **Billabongs** Highly Voted 1 year, 10 months ago

Looks Correct

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit>

<https://docs.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if?tabs=azure-powershell>

upvoted 18 times

 **surprise0011** 1 year, 2 months ago

received 2023-04-17 went given answer, score 926

upvoted 11 times

 **Esward** 1 year, 5 months ago



Agreed

upvoted 1 times

 **drhuu** Highly Voted 1 year, 8 months ago

on exam oct 2022

upvoted 6 times

  **seijumi** 1 year, 8 months ago

Whats ur score?

upvoted 1 times

  **MikeAWS** Most Recent 9 months ago

The Azure Resource Manager template (ARM template) test toolkit checks whether your template uses recommended practices.

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/test-toolkit>

Before deploying an Azure Resource Manager template (ARM template), you can preview the changes that will happen. Azure Resource Manager provides the what-if operation to let you see how resources will change if you deploy the template. The what-if operation doesn't make any changes to existing resources. Instead, it predicts the changes if the specified template is deployed.

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/templates/deploy-what-if?tabs=azure-powershell>

upvoted 3 times

  **longnguyendh** 11 months ago

Correct

upvoted 1 times

  **Saluk\_DE** 1 year, 3 months ago

Question was in Exam 2023-03-30

upvoted 1 times

  **sarmaria** 1 year, 3 months ago



Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 1 times

  **mabdo** 1 year, 4 months ago

Was om exam 02/2023


upvoted 1 times

  **kociuba** 1 year, 6 months ago

got it on my exam 30-12-2022


score: 818

upvoted 4 times

  **OPT\_001122** 1 year, 8 months ago

is the given answer correct?

upvoted 1 times

  **Parkruta** 1 year, 8 months ago

on exam 10/28/2022

upvoted 3 times

  **vcfvct** 1 year, 8 months ago

Got this on 10/21/2022., score 850

upvoted 3 times

  **gmishra88** 1 year, 8 months ago

Don't forget which one falls into which one. It can be confusing

upvoted 2 times

  **coffecold** 1 year, 8 months ago

Correct

<https://4bes.nl/2020/08/09/testing-arm-templates/>

upvoted 2 times

You develop Azure Durable Functions to manage vehicle loans.

The loan process includes multiple actions that must be run in a specified order. One of the actions includes a customer credit check process, which may require multiple days to process.

You need to implement Azure Durable Functions for the loan process.

Which Azure Durable Functions type should you use?

- A. orchestrator
- B. client
- C. entity
- D. activity

**Suggested Answer: A**

Community vote distribution

A (100%)

🗳️ **imanonion** Highly Voted 1 year, 11 months ago

**Selected Answer: A**

An orchestrator function is the appropriate type of Azure Durable Function to use in this scenario, because it allows you to define the overall flow of the loan process and call other functions or activities as needed. The credit check process can be implemented as a separate activity function, which can be called by the orchestrator function and run in parallel with other actions in the loan process.

Entity functions are designed for use cases where you need to perform operations on a shared piece of state in a reliable and atomic way, such as a distributed queue or counter. In this scenario, it does not appear that there is a need to use entity functions.

There is no such thing as a "client" function in Azure Durable Functions. "Activity" functions are called by orchestrator functions to perform specific tasks, but the orchestrator function is the one that defines the overall flow of the loan process, so it is the correct answer in this case.

upvoted 27 times

🗳️ **oskx2** 10 months ago

There are client functions, read the doc: <https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-types-features-overview#client-functions>

"The primary way to deliver these messages is by using an orchestrator client binding or an entity client binding from within a client function."

upvoted 6 times

🗳️ **surprise0011** 1 year, 8 months ago

received 2023-04-17 went given answer, score 926

upvoted 3 times

🗳️ **ffkraemer** Most Recent 1 year ago

**Selected Answer: A**

The key is "specified order" so A. Orchestrator

upvoted 4 times

🗳️ **Saluk\_DE** 1 year, 9 months ago

Question was in Exam 2023-03-30

upvoted 1 times

🗳️ **sarmaria** 1 year, 9 months ago



Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 2 times

🗳️ **NombreFalso** 1 year, 10 months ago

**Selected Answer: A**

If we can only choose one then A, would also include D if possible  
upvoted 3 times

  **Esward** 1 year, 11 months ago

Orchestrator

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-sequence?tabs=csharp>

upvoted 2 times

## HOTSPOT

-

You are developing an Azure Function app.

All functions in the app meet the following requirements:

- Run until either a successful run or until 10 run attempts occur.
- Ensure that there are at least 20 seconds between attempts for up to 15 minutes.

You need to configure the host.json file.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```
{
  "retry": {
    "strategy": "exponentialBackoff",
    "maxRetryCount": 10,
    "healthCheckInterval": "00:00:20",
    "healthCheckThreshold": "00:15:00"
  }
}
```


## Answer Area

```
{
  "retry": {
    "healthMonitor": "singleton",
    "strategy": "exponentialBackoff",
    "counterThreshold": 10,
    "maxRetryCount": 10,
    "healthCheckInterval": "00:00:20",
    "healthCheckThreshold": "00:15:00"
  }
}
```

Suggested Answer:

 **uffuchsi** Highly Voted 2 years, 4 months ago

Received this in my exam today (22/02/2023). Selected retry, exponentialBackoff, and maxRetry. Score 927.  
upvoted 19 times

 **imanonion** Highly Voted 2 years, 5 months ago

Answer seems correct.

Did some research about exponential backoff:

The "exponential backoff" retry strategy is a technique for retrying failed operations in a manner that avoids overloading the system being accessed. It works by increasing the amount of time that is waited between each retry attempt, using an exponential function to calculate the wait time.

For example, with a coefficient of 2.0 and an initial retry interval of 1 second, the wait times between retries might be 1 second, 2 seconds, 4 seconds, 8 seconds, and so on. This allows the system being accessed to recover from any failures or load spikes before the next retry attempt is made, reducing the likelihood of further failures.

upvoted 18 times

 **the\_odd\_one** 1 year, 8 months ago


Where is then the coefficient argument in the code? or is this an implicit value?

upvoted 2 times

 **dom271219** 1 year, 9 months ago

So, based on your explanations, exponentialBackoff may conflict with the minimumInterval (20s) and maximumInterval (15m).

upvoted 1 times

 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

 **jd99** 10 months, 2 weeks ago

For exponential backoff min/max interval are used. For fixed delay delayInterval is used. My answer would be retry - exponentialBackoff - maxRetryCount

<https://learn.microsoft.com/en-us/java/api/com.microsoft.azure.functions.annotation.exponentialbackoffretry?view=azure-java-archive>

<https://learn.microsoft.com/en-us/java/api/com.microsoft.azure.functions.annotation.fixeddelayretry?view=azure-java-archive>

upvoted 2 times

 **At09** 1 year ago

retry, exponentialBackoff, and maxRetry

upvoted 1 times

🗨️ 👤 **neelkanths** 1 year, 2 months ago

Got it on 20 April 2024...Marks > 950...answer is correct....all questions from examtopics 400 question bank

Answer -

retry

exponential backoff

maxRetry

Exponential doesn't interfere with maximum time, it calculates time interval based on maximum time.

upvoted 3 times

🗨️ 👤 **Kk46** 1 year, 1 month ago

Do you have all the questions

upvoted 2 times

🗨️ 👤 **evandrorenan** 1 year, 4 months ago

Fixed delay.

With exponentialBackoff the number of retries would be less than 20 in 15 minutes.

CounterThreshold would need a delay field. Since it's not provided, it would go with the default mechanism, which is also exponential backoff.

upvoted 2 times

🗨️ 👤 **1989jay** 1 year, 5 months ago

This question was on exam. 09/01/2024. Passed 872.

upvoted 1 times

🗨️ 👤 **bgbgvfvf** 1 year, 6 months ago

I think given answer is correct.

upvoted 1 times

🗨️ 👤 **85e4c91** 1 year, 8 months ago

I don't get it, Exponential Backoff would increase the time between retries, so it would not be 20 seconds anymore and it is clearly stated that the time between retries should be 20 seconds for the 15 minutes. Both Fixed delay and Exponential backoff have Max retry count.

upvoted 2 times

🗨️ 👤 **asdfdude** 1 year, 2 months ago

It says "at least 20 seconds"

upvoted 1 times

🗨️ 👤 **85e4c91** 1 year, 8 months ago

Per documentation on <https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=exponential-backoff%2Cisolated-process&pivots=programming-language-csharp>

Fixed Delay: A specified amount of time is allowed to elapse between each retry.

Exponential Backoff: The first retry waits for the minimum delay. On subsequent retries, time is added exponentially to the initial duration for each retry, until the maximum delay is reached. Exponential back-off adds some small randomization to delays to stagger retries in high-throughput scenarios.

upvoted 2 times

🗨️ 👤 **leviatas** 1 year, 9 months ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 1 times

🗨️ 👤 **p2006** 1 year, 9 months ago

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=exponential-backoff%2Ccsharp-script&pivots=programming-language-csharp#retry-examples>

upvoted 1 times

🗨️ 👤 **p2006** 1 year, 9 months ago

"Ensure that there are at least 20 seconds between attempts for up to 15 minutes." ==> exponentialBackoff

upvoted 1 times

🗨️ 👤 **longnguyendh** 1 year, 11 months ago



the answer is correct

retry

exponentialBackoff

maxRetryCount

upvoted 1 times

🗒️ 👤 **NightshadeRC** 1 year, 11 months ago

Had this question today: 2023-07-26

upvoted 3 times

🗒️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 3 times

🗒️ 👤 **PP2015** 2 years, 3 months ago

Received this in my exam today (29/02/2023). Selected retry, exponentialBackoff, and maxRetry. Score 850.

upvoted 3 times

🗒️ 👤 **mmdex** 2 years, 4 months ago

retry

exponentialBackoff

maxRetryCount

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages?tabs=exponential-backoff%2Ccsharp-script&pivots=programming-language-csharp#retry-strategies>

Only fixedDelay and exponentialBackoff strategies are possible and fixedDelay has different parameters than shown here.

upvoted 4 times

You develop Azure Web Apps for a commercial diving company. Regulations require that all divers fill out a health questionnaire every 15 days after each diving job starts.

You need to configure the Azure Web Apps so that the instance count scales up when divers are filling out the questionnaire and scales down after they are complete.

You need to configure autoscaling.


What are two possible auto scaling configurations to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Recurrence profile
- B. CPU usage-based autoscaling
- C. Fixed date profile
- D. Predictive autoscaling

**Suggested Answer:** CD

Community vote distribution



**Osip21** Highly Voted 2 years, 5 months ago

**Selected Answer:** BD

I think it should

- B. CPU usage-based autoscaling
- D. Predictive autoscaling

A. Recurrence profile is used to schedule the scaling of resources at specific times or dates, but it does not meet the requirement to scale up when divers are filling out the questionnaire and scale down after they are complete. It only triggers scaling based on a set schedule, not based on actual usage.

C. Fixed date profile is used to specify the number of instances at a specific date and time, but it also does not meet the requirement to dynamically scale based on actual usage. It only sets a fixed number of instances and does not adjust based on changing workloads.

upvoted 37 times

**RikinPatel** 1 year, 9 months ago

Predictive autoscale - Support is only available for virtual machine scale sets.

D is incorrect options

A - Recurrence profile is correct because you can specify every 15 days.

upvoted 3 times

**antonmelode** 1 year ago

Recurrence profile occurs weekly only

upvoted 4 times

**Shayman204** 1 year, 8 months ago

You have many drivers and they dont always fall on same 15day cycle? correct? I think B and D is right also because they mentioned "scaled UP" not scale "OUT" which also validates further its Azure web apps on a VM

upvoted 7 times

**imanonion** Highly Voted 2 years, 5 months ago

**Selected Answer:** AC

I think it should be A and C instead because:

Options "B. CPU usage-based autoscaling" and "D. Predictive autoscaling" are not relevant to this scenario, because they are based on resource utilization or predictive analytics rather than a fixed schedule.

upvoted 13 times

  **TheFivePips** 10 months, 1 week ago

The question should be more clear, but they seem to be saying that each individual driver needs to fill out the form 15 days after a job begins, not that everyone will be filling out the form every 15 days. Just a poorly worded question making this needlessly confusing



upvoted 1 times

  **Test123\_456** Most Recent 4 days, 3 hours ago

Selected Answer: AB

It should be AB

upvoted 1 times

  **2f17d52** 2 months ago

Selected Answer: AB

From Chatgpt:

- A. Recurrence profile | ✓ Good | You can schedule scaling based on a recurring time pattern (example: scale up every 15th day or every 2nd Friday).
- B. CPU usage-based autoscaling | ✓ Good | If more divers fill in forms, CPU load rises → autoscaling on CPU makes sense.
- C. Fixed date profile | ✗ Not ideal | A fixed date profile is for one-time events (like a Black Friday sale, not every 15 days).
- D. Predictive autoscaling | ✗ Not the best fit | Predictive autoscaling uses machine learning to predict patterns – but it needs historical data and is mainly useful for very regular patterns (like daily). It's not precise for 15 days after random dive jobs.

upvoted 1 times

  **Thameur01** 3 months ago

Selected Answer: AC

✓ A. Recurrence profile

This is valid because the questionnaire needs to be filled every 15 days after each job starts, which follows a recurring schedule. You can configure autoscaling based on a recurring time-based rule to scale up during expected usage (e.g., every 15 days for a few hours) and scale down afterward.

✓ C. Fixed date profile

This is also valid if you know the exact dates when scaling should occur (e.g., a known diving schedule). A fixed date profile allows scaling rules to be triggered based on specific calendar dates and times, perfect for predictable, one-time or irregular scheduled events like diving jobs.

upvoted 1 times

  **borsooq** 4 months, 2 weeks ago

Selected Answer: BC

First of all we have App Service here, not a Scale Set, so - predictive autoscaling is not available. If there is anyone that would show me where to set it - I would appreciate it and probably change my mind.

Since B is obvious and is definitely one of the solutions, the only ones that left are A and C.

Recurrence - I guess it is Repeat specific days profile, because I can't see any other where I could define any recurrence. Can't see possibility to set Recurrence profile to - every 15 days as it is per day of week. So the only way to do it would be fixed days - simply specific days in the year on which you have to set number of instances. That's it, so my answer would be B and C.

upvoted 1 times

  **murat12345** 9 months ago

A and C.

I tried auto-scaling a web-app, and both "Fixed date profiles" and "Recurring profiles" are possible options.

"CPU usage-based autoscaling" can be configured as part of "Fixed date profiles" and "Recurring profiles".

<https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-multiprofile?tabs=templates>

<https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-get-started?toc=%2Fazure%2Fapp-service%2Ftoc.json>

upvoted 1 times

  **freudn** 1 year ago

Selected Answer: BC

Long winded but bare with me:

upvoted 1 times

  **anishrawal** 1 year, 3 months ago

It should only be B.

Not recurrence as it scales on a particular date and time or day of the week.

Not Predictive as it uses ML to forecast based on historical CPU usage patterns.

Not Fixed as well.

<https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-predictive>

upvoted 1 times

🗨️ **Chriscrea8** 1 year, 3 months ago

To configure autoscaling for the Azure Web Apps to scale up when divers are filling out the questionnaire and scale down after they are complete, you can use the following auto scaling configurations:

A. Recurrence profile - You can set up a recurring schedule to scale out the instances during the times when divers are expected to fill out the questionnaire (e.g., every 15 days after a diving job starts) and scale in when the questionnaire completion period ends.

B. CPU usage-based autoscaling - You can configure autoscaling based on CPU usage metrics. When the number of divers filling out the questionnaire increases, it might lead to increased CPU usage, triggering the autoscaling to add more instances. Similarly, when the questionnaire completion period ends, CPU usage might decrease, leading to scaling down the instances.

So, the correct answers are A (Recurrence profile) and B (CPU usage-based autoscaling).

upvoted 1 times

🗨️ **raymond\_abcd** 1 year, 4 months ago

**Selected Answer: BD**

See: <https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-understanding-settings#autoscale-profiles>

It is not the fixed-date profile because it is relevant for a single date and time. The profile runs only once

The Recurrence profile: A recurrence profile is used for a day or set of days of the week. The schema for a recurring profile doesn't include an end date.

upvoted 3 times

🗨️ **raymond\_abcd** 1 year, 4 months ago

Correction it is A and B

upvoted 1 times

🗨️ **littleprince** 1 year, 4 months ago

**Selected Answer: BD**

Predictive autoscale uses machine learning to help manage and scale Azure Virtual Machine Scale Sets with cyclical workload patterns. It forecasts the overall CPU load to your virtual machine scale set, based on your historical CPU usage patterns. It predicts the overall CPU load by observing and learning from historical usage. This process ensures that scale-out occurs in time to meet the demand.

Predictive autoscale needs a minimum of 7 days of history to provide predictions. The most accurate results come from 15 days of historical data.

upvoted 2 times

🗨️ **manopeydakon** 1 year, 5 months ago

A and B,

For autoscaling in the context of Azure Web Apps with the described requirements, the best options would be:

B. CPU usage-based autoscaling

This would allow you to scale up the instance count when there is increased load during the questionnaire filling periods.

A. Recurrence profile

You can schedule the autoscaling based on a recurrence profile to handle the periodic nature of the health questionnaire every 15 days.

These two options should provide a suitable configuration for scaling up during questionnaire periods and scaling down afterward.

upvoted 4 times

🗨️ **alirasouli** 1 year, 7 months ago

**Selected Answer: AB**

B. CPU usage-based is a common practice for autoscaling.

A. Profile is for date and time-related autoscaling. Do not overthink! Every 15 days => recurring profile

Why C & D are wrong?

C. Simply do not overthink!

D. Predictive autoscaling is for VMSS, not Azure Web Apps.

upvoted 2 times

🗨️ 👤 **ami2023** 1 year, 7 months ago

**Selected Answer: BD**

[15 days after 'each' diving job starts] which means the day to fill the questionnaires can fall in any day, therefore the options that are scaling based on time/dates (A, C) cannot be the solution.

upvoted 3 times

🗨️ 👤 **Ciupaz** 1 year, 7 months ago

Hope don't have this question in my exam.

upvoted 5 times

🗨️ 👤 **JoseManel** 1 year, 7 months ago

did you get this in your exam?

upvoted 1 times

🗨️ 👤 **Ciupaz** 1 year, 7 months ago

I haven't taken the exam yet.

upvoted 1 times

🗨️ 👤 **Teit\_examtopics** 1 year, 9 months ago

<https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-predictive>

Only applies to VM Scale sets

upvoted 2 times

## HOTSPOT

-


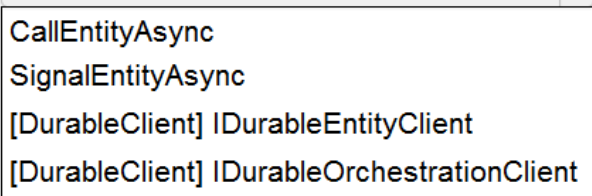
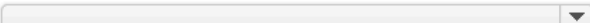
You are developing an online game that allows players to vote for their favorite photo that illustrates a word. The game is built by using Azure Functions and uses durable entities to track the vote count.

The voting window is 30 seconds. You must minimize latency.

You need to implement the Azure Function for voting.

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

**Answer Area**

```
[FunctionName("Vote")]
public static async Task<HttpResponseMessage> Run(
    [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
     c,
    
    string id)
{
    var eid = new EntityId("pic", id);
    await c.  (eid, "vote");

    return req.CreateResponse(HttpStatusCode.OK);
}
```

The image shows a code editor with a C# function for voting. There are two dropdown menus for selecting the correct completion for the `c` parameter in the `Run` method. The first dropdown menu is open, showing the following options: `CallEntityAsync`, `SignalEntityAsync`, `[DurableClient] IDurableEntityClient`, and `[DurableClient] IDurableOrchestrationClient`. The second dropdown menu is also open, showing the same four options. The correct completion for the `c` parameter is `[DurableClient] IDurableEntityClient`.

### Answer Area

```
[FunctionName("Vote")]
public static async Task<HttpResponseMessage> Run(
    [HttpTrigger("POST", Route = "pic/{id}")] HttpRequestMessage req,
    c,
    string id)
{
    var eid = new EntityId("pic", id);
    await c.
    return req.CreateResponse(HttpStatusCode.OK);
}
```

CallEntityAsync  
SignalEntityAsync  
[DurableClient] IDurableEntityClient  
[DurableClient] IDurableOrchestrationClient

CallEntityAsync  
SignalEntityAsync  
[DurableClient] IDurableEntityClient  
[DurableClient] IDurableOrchestrationClient

Suggested Answer:

 **d79** Highly Voted 1 year, 8 months ago

IDurableEntityClient is correct to works with entities.

Signal is used for 1-way

Call would be used for 2-way

upvoted 9 times

 **12jan37** Highly Voted 1 year, 8 months ago

Answer is correct: <https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-dotnet-entities>

upvoted 9 times


 **130nk3r5** 12 months ago

Got this today.

Went with answer here.

Score 927

upvoted 6 times

 **Ciupaz** 11 months, 4 weeks ago

Good score!

upvoted 1 times

 **Teit\_examptopics** Most Recent 1 year, 3 months ago

<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-dotnet-entities#example-client-signals-entity>

upvoted 1 times

 **p2006** 1 year, 3 months ago


<https://learn.microsoft.com/en-us/azure/azure-functions/durable/durable-functions-dotnet-entities#example-client-signals-entity>

upvoted 1 times

 **adilkhan** 1 year, 8 months ago


Answers are correct. Also Chat GPT generated the same code.

upvoted 2 times

 **notedo** 1 year, 8 months ago

This seems to be a new question. It's not here last month. Will exam next week

upvoted 3 times

 **pelsky** 1 year, 8 months ago

Same, I'm writing next week

upvoted 1 times

HOTSPOT

-

You have an App Service plan named asp1 based on the Free pricing tier.

You plan to use asp1 to implement an Azure Function app with a queue trigger. Your solution must minimize cost.

You need to identify the configuration options that will meet the requirements.

Which value should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

#### Configuration option

#### Value

Azure App Service feature

Always On  
Managed identity  
Continuous deployment

Azure App Service pricing tier

Basic  
Shared  
Standard

### Answer Area

#### Configuration option

#### Value

Azure App Service feature

Always On  
**Managed identity**  
Continuous deployment

Azure App Service pricing tier

**Basic**  
Shared  
Standard

Suggested Answer:

 **[Removed]**  1 year, 11 months ago

1. Always on

If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. On an App Service plan, the functions runtime goes idle after a few minutes of inactivity, so only HTTP triggers will "wake up" your functions. The Always on setting is available only on an App Service plan. On a Consumption plan, the platform activates function apps automatically.

Even with Always On enabled, the execution timeout for individual functions is controlled by the functionTimeout setting in the host.json project file.

2. Basic

The dedicated App Service plans supported for function app hosting include Basic, Standard, Premium, and Isolated SKUs. Free and Shared tier App Service plans aren't supported by Azure Functions. Basic is the correct answer since it is cheaper than Standard.



<https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan>

upvoted 28 times

🗨️ 👤 **1CY1** 1 year ago

While I will run with this answer it could be argued that a Shared Plan minimizes costs and you can run the function app on a consumption plan.

upvoted 1 times

🗨️ 👤 **RikinPatel** 1 year, 9 months ago

In your ans. you are talking about HTTP trigger but as per que. requirement is "queue trigger"

upvoted 1 times

🗨️ 👤 **Gauravbio** 1 year, 2 months ago

We need Always on: otherwise the function app is not recognised, and only Basic and above plans supports function apps.

upvoted 1 times

🗨️ 👤 **alirasouli** Highly Voted 👍 1 year, 7 months ago

Always On and Continuous Deployment are not supported on the Free tier plans.

Answer: Managed Identity & Basic

upvoted 9 times

🗨️ 👤 **Thameur01** Most Recent 🕒 3 months ago

✓ Managed identity

Queue-triggered Azure Functions do not require Always On if hosted in the Consumption plan.

Since your plan is asp1 (Free), the most applicable feature from the dropdown that aligns with cost efficiency and functionality is Managed identity (used for secure access to Azure services, e.g., queues).

Always On is not supported in Free tier and is not necessary for queue triggers in Consumption plan.

Continuous deployment is unrelated to queue triggering or minimizing cost.

✓ Shared or Basic

Free and Shared tiers are the lowest-cost options.

However, Azure Function Apps with queue triggers typically require the Consumption plan, which is not shown in this dropdown.

Among the available options, Shared is the best low-cost tier that supports more functionality than Free (though the real best choice in production would be Consumption, which is not listed here).

Standard supports Always On and more, but is more costly.

upvoted 1 times

🗨️ 👤 **Dkijc** 8 months, 1 week ago

I believe it's

1. Managed Identity

2. Basic

Source: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-identity-based-connections-tutorial-2>

upvoted 1 times

🗨️ 👤 **fiercelfreedom** 8 months, 2 weeks ago

Use managed identity to connect to the queue that's triggering the function execution. Azure Functions on a consumption plan don't require "Always On" for auto-scaling. The "Always On" feature is irrelevant here.

upvoted 2 times

🗨️ 👤 **fuzzzzel** 1 year, 2 months ago

Regarding this article a connection to a queue can be identity based and Azure functions uses a managed identity here. Maybe that is why the answer states to use this feature.

<https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan>

upvoted 2 times

🗨️ 👤 **antonmelode** 1 year ago

That's not really an app service feature though. It's more an Azure Authentication feature

upvoted 1 times

🗨️ 👤 **Will\_I\_Am\_586** 1 year, 3 months ago

I didn't get it. Cos in the Microsoft docs says "Free and Shared tier App Service plans aren't supported by Azure Functions...", So... this kind of config for the App Service Plan should work?

<https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan>

upvoted 1 times

🗳️ 👤 **LJNZ2049** 1 year, 6 months ago

Always on: "If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. "

Basic: "Free and Shared tier App Service plans aren't supported by Azure Functions. For a lower-cost option hosting your function executions, you should instead consider the Consumption plan, where you are billed based on function executions."

upvoted 2 times

🗳️ 👤 **Sakhana88** 1 year, 6 months ago

1. Always On

2. Basic

upvoted 2 times

🗳️ 👤 **lord\_okello** 1 year, 8 months ago

Answer is correct. <https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan>

upvoted 2 times

🗳️ 👤 **Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 3 times

🗳️ 👤 **leviatas** 1 year, 9 months ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 2 times

🗳️ 👤 **longnguyendh** 1 year, 11 months ago

1. Always on

2. Basic

upvoted 3 times

🗳️ 👤 **AlexMMReis** 1 year, 11 months ago

1. Always on

2. Basic

"If you run on an App Service plan, you should enable the Always on setting so that your function app runs correctly. On an App Service plan, the functions runtime goes idle after a few minutes of inactivity, so only HTTP triggers will "wake up" your functions. The Always on setting is available only on an App Service plan."

<https://learn.microsoft.com/en-us/azure/azure-functions/dedicated-plan>

upvoted 2 times

🗳️ 👤 **AlexMMReis** 1 year, 11 months ago

Regarding why basic:

Is the the cheapest tier with dedicated compute type

upvoted 1 times

🗳️ 👤 **ayush66203** 1 year, 8 months ago

but there is not any requirement of having a dedicated computation? i think shared would shared pricing tier will be good here.

upvoted 1 times

🗳️ 👤 **75daltonic** 1 year, 11 months ago

1. Always on

2. Free

upvoted 1 times

🗳️ 👤 **75daltonic** 1 year, 11 months ago

Sorry, 2. Basic :-)

upvoted 2 times

🗳️ 👤 **Roel1** 1 year, 11 months ago

Managed identity?

upvoted 4 times

DRAG DROP

-

You are developing several microservices to run on Azure Container Apps.

The microservices must allow HTTPS access by using a custom domain.

You need to configure the custom domain in Azure Container Apps.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

### Actions

Validate the custom domain name.

Enable ingress.

Bind the certificate.

Add DNS records to the domain provider.

Add the custom domain name.

### Answer Area



### Answer Area

Enable ingress.

Add the custom domain name.

Validate the custom domain name.

Bind the certificate.

Add DNS records to the domain provider.

Suggested Answer:

**andrewnm** Highly Voted 1 year, 11 months ago

from my understanding of <https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates>

- 1) Enable Ingress == step 2
- 2) Add the custom domain name == step 4
- 3) Bind certificate == steps 5 to 11
- 4) Add DNS records to the domain provider == 12 to 14
- 5) Validate the custom domain name == step 15

Using similar questions as a basis, I'm assuming that "Add the custom.." means click the button and not step 16 "Once validation succeeds, select the Add button"

Not very clear.

upvoted 45 times

**fuji36850** 1 year, 1 month ago

You need to have DNS records first to validate if its working

upvoted 3 times

**75daltonic** 1 year, 11 months ago

This question is \$%&#\*!!!

If you follow the steps in <https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates>, you have to validate the domain name twice.

In some videos, as @Desniskam has posted, you can see that first of all you fill the domain name and after it, you click on validate button. Then,

other configuration options appears. After filling all other configurations, you have to click on validate button again.

My doubt is that I have watched some videos and it is different for app services, web app, etc. I haven't found a container example, so I will follow microsoft learn steps, with validating as last step.

upvoted 1 times

  **andrewnm** 1 year, 11 months ago

Occurs to me that steps 5-11 adding the cert maybe irrelevant and binding the cert is actually the last step, based on what you click in the UI.. so I'd revise that to:

- 1) Enable Ingress
- 2) Add the custom domain name
- 3) Add DNS records to the domain provider
- 4) Validate the custom domain name
- 5) Bind certificate,

Still not clear!

upvoted 7 times

  **Papee** Highly Voted 1 year, 11 months ago



1. Enable ingress: Enable ingress for the Azure Container Apps. This allows external traffic to reach the microservices.
2. Add the custom domain name: Add the custom domain name that you want to use for HTTPS access to your Azure Container Apps.
3. Add DNS records to the domain provider: After adding the custom domain name in Azure Container Apps, you need to add the required DNS records (such as CNAME or A records) to your domain provider's DNS settings. This step is essential for directing traffic from the custom domain to your Azure Container Apps.
4. Validate the custom domain name (after adding the dns records to the provider): After adding the DNS records to the domain provider, you need to validate the custom domain name in Azure Container Apps. This step ensures that the DNS records are correctly configured and the domain is pointing to your Azure Container Apps.
5. Bind the certificate (only enabled after finishing the validation): Once the custom domain name is validated, you can bind the SSL/TLS certificate to enable HTTPS access for the custom domain.

upvoted 33 times

  **neelkanths** 1 year, 2 months ago

how come this answer comes - while clearly it mentions certificate is added before adding dns records?

upvoted 1 times

  **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 2 times

  **Thameur01** 3 months ago

Enable ingress

- Enables external HTTP/HTTPS traffic to your Container App.

Add the custom domain name

- Specify the custom domain in your Azure Container App settings.

Add DNS records to the domain provider

- Add CNAME or A records to your DNS provider so that Azure can validate the domain.



Validate the custom domain name

- Azure verifies that the domain is correctly pointed to your app.

Bind the certificate

- Attach an SSL certificate to enable HTTPS access.

upvoted 1 times

  **VithobaS** 5 months, 1 week ago

1. Ingress must be enabled first to allow external access.
2. Adding the custom domain to the Azure Container App is required before mapping DNS.
3. DNS records must be updated at the domain provider to point the domain to the app.

4. Validation ensures the DNS changes have propagated and the domain is correctly linked.

5. Finally, binding a certificate enables HTTPS for secure access.

upvoted 1 times

🗨️ 👤 **neelkanths** 1 year, 2 months ago

Per copilot -

To configure a custom domain in Azure Container Apps and allow HTTPS access, you should perform the actions in the following order:

1) Enable Ingress: This will allow external access to your microservices.

2) Add custom domain name: You need to specify the custom domain that you want to use for your microservices.

3) Bind the certificate: This is necessary for enabling HTTPS. You need to bind the SSL/TLS certificate to your custom domain.

4) Add DNS Records to the domain provider: This will point your custom domain to the IP address of your Azure Container Apps.

5) Validate the custom domain name: This is the final step to ensure that your custom domain has been properly configured and is accessible.

upvoted 3 times

🗨️ 👤 **neelkanths** 1 year, 2 months ago

given answer is correct

upvoted 1 times

🗨️ 👤 **manopeydaikon** 1 year, 5 months ago

Navigate to your container app in the Azure portal

Verify that your app has ingress enabled by selecting Ingress in the Settings section. If ingress is not enabled, enable it with these steps:

Set HTTP Ingress to Enabled. = 1

Select the desired Ingress traffic setting.

Enter the Target port.

Select Save.

Under the Settings section, select Custom domains.

Select the Add custom domain button.

In the Add custom domain and certificate window, in TLS/SSL certificate, select Bring your own certificate. = 2

In domain, enter the domain you want to add.

Select Add a certificate. = 3

In the Add certificate window, in Certificate name, enter a name for this certificate.

In Certificate file section, browse for the certificate file you want to upload.

Select Validate. = 4

Once validation succeeds, select Add.

In the Add custom domain and certificate window, in Certificate, select the certificate you just added. = 5

Select the Hostname record type based on the type of your domain.

upvoted 2 times

🗨️ 👤 **AzDeveloper** 1 year, 6 months ago

1. Enable ingress.

2. Add the custom domain name.

3. Bind the certificate.

4. Add DNS records to the domain provider.

5. Validate the custom domain name.

According to "Custom domain names and bring your own certificates in Azure Container Apps" article.

<https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates>

upvoted 4 times

🗨️ 👤 **Weam** 1 year, 6 months ago

another useful link with screenshots: <https://www.examttopics.com/exams/microsoft/az-204/view/9/>

Answer:

1) enable ingress

2) Add custom domain name

3) bind certificate

4) Add DNS records to doman provider

5) validate

upvoted 4 times


🗨️ 👤 **dddddd111** 1 year, 7 months ago

Please look at No.2, No. 4, No.7, No. 10, No.14 --> This is the correct order. Note: No. 15 is validating the DNS not the CDN.

Correct answer.

1. Enable Ingress
2. Add the custom domain name
3. Bind certificate
4. Validate the custom domain name
5. Add DNS records to the domain provider

upvoted 11 times

 **Vukile** 1 year, 7 months ago

On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors

upvoted 3 times

 **Venusmappi** 1 year, 8 months ago

- 1) Enable Ingress
- 2) Add the custom domain name
- 3) Add DNS records to the domain provider
- 4) Validate the custom domain name
- 5) Bind certificate

upvoted 2 times

 **HimGarg** 1 year, 9 months ago

Why fixed order is required. You can do it in different ways.

upvoted 2 times

 **NicolasMous** 1 year, 9 months ago

The correct answer is:

1. Enable ingress
2. Add the custom domain name
3. Bind the certificate
4. Add DNS records to the domain provider
5. Validate the custom domain name

Explanation:

-Enable ingress: This is the first step as it allows HTTP access to your container app.

-Add the custom domain name: After enabling ingress, you can add the custom domain name to your container app.

-Bind the certificate: Once the custom domain name is added, you can bind the certificate to it. This can be a certificate you upload or a free managed certificate provided by Azure.

-Add DNS records to the domain provider: After binding the certificate, you need to add DNS records to your domain provider. This points the domain to your container app and verifies that you own it.


-Validate the custom domain name: Finally, you validate the custom domain name. This ensures that the DNS records have been correctly configured and that the domain correctly points to your container app.

upvoted 7 times

 **Hasti123** 1 year, 9 months ago

100% right

upvoted 1 times

 **Kanasan** 1 year, 10 months ago

1. Enable ingress
2. Add the custom domain name
3. Bind the certificate
4. Add DNS records to the domain provider
5. Validate the custom domain name

<https://learn.microsoft.com/en-us/azure/container-apps/custom-domains-certificates>

First validation is for certification, not for a custom domain name.

So, I put the validation step as the last one, because we validate the custom domain at the end.

upvoted 3 times

  **longnguyendh** 1 year, 11 months ago

<https://learn.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-custom-domain?tabs=root%2Cazurecli>

I think that It should be

- 1) Add the custom domain name
- 2) Bind certificate
- 3) Validate the custom domain name
- 4) Add DNS records to the domain provider
- 5) Enable Ingress

upvoted 2 times

  **1CY1** 1 year ago

Enable Ingress must be first to allow access to the container.

upvoted 1 times

You are developing several microservices to run on Azure Container Apps. External HTTP ingress traffic has been enabled for the microservices.

The microservices must be deployed to the same virtual network and write logs to the same Log Analytics workspace.

You need to deploy the microservices.

What should you do?

- A. Enable single revision mode.
- B. Use a separate environment for each container.
- C. Use a private container registry image and single image for all containers.
- D. Use a single environment for all containers.
- E. Enable multiple revision mode.

**Suggested Answer: A**

Community vote distribution

D (100%)

🗳️ 👤 **MikeM27** Highly Voted 1 year, 5 months ago

The correct answer should be D  
upvoted 12 times

🗳️ 👤 **Steffexx** 1 year, 5 months ago

That is what I thought as well. The single revision would not solve anything here.  
upvoted 3 times

🗳️ 👤 **Regex37** 5 months ago

A bit more background: "Use a single environment when you want to: [...] Have applications share the same log destination"  
(<https://learn.microsoft.com/en-us/azure/container-apps/environment>)  
upvoted 1 times

🗳️ 👤 **J\_obando13** Most Recent 1 month, 2 weeks ago

**Selected Answer: D**

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)  
upvoted 1 times

🗳️ 👤 **Sushantjadhav416** 9 months, 1 week ago

**Selected Answer: D**

<https://www.examtips.com/exams/microsoft/az-204/view/9/#>  
upvoted 2 times

🗳️ 👤 **jobolesonihal** 9 months, 1 week ago

Per copilot -

To configure a custom domain in Azure Container Apps and allow HTTPS access, you should perform the actions in the following order:

- 1) Enable Ingress: This will allow external access to your microservices.
  - 2) Add custom domain name: You need to specify the custom domain that you want to use for your microservices.
  - 3) Bind the certificate: This is necessary for enabling HTTPS. You need to bind the SSL/TLS certificate to your custom domain.
  - 4) Add DNS Records to the domain provider: This will point your custom domain to the IP address of your Azure Container Apps.
  - 5) Validate the custom domain name: This is the final step to ensure that your custom domain has been properly configured and is accessible.
- upvoted 2 times

🗳️ 👤 **Xtoluck** 10 months, 2 weeks ago

The correct answer should be D  
upvoted 1 times



🗳️ 👤 **Rabi3** 11 months, 2 weeks ago

**Selected Answer: D**

Answer should be D

upvoted 1 times

🗳️ 👤 **SanjanaBAZINGAA** 1 year ago

D. Use a single environment for all containers.

Explanation:

Using a single environment for all containers can help in managing multiple microservices within the same context, which may include deploying them to the same virtual network and configuring them to write logs to the same Log Analytics workspace.

The concept of "environment" in this context could refer to a shared Azure Container Apps instance or a broader resource group or Azure region where the microservices are deployed and managed together.<sup>4</sup>

upvoted 3 times

🗳️ 👤 **dhivya6613** 1 year ago

D. Use a single environment for all containers: Azure Container Apps Environment acts as a logical grouping of container apps that share common characteristics. By using a single environment for all your microservices, you can easily manage network and security settings across all apps, ensuring they are deployed in the same virtual network and can write logs to the same Log Analytics workspace. This approach simplifies network and log management while maintaining isolation at the microservice level.

upvoted 3 times

🗳️ 👤 **Karunan** 1 year, 1 month ago

If D is the correct answer, why answer is not changed?

upvoted 1 times

🗳️ 👤 **CarlosTheBoldest** 1 year, 1 month ago

**Selected Answer: D**

D FTW!

upvoted 2 times

🗳️ 👤 **p2006** 1 year, 3 months ago

**Selected Answer: D**

<https://learn.microsoft.com/en-us/azure/container-apps/environment#single-environment>

upvoted 4 times

🗳️ 👤 **renierdbuyn** 1 year, 4 months ago

**Selected Answer: D**

D is correct

upvoted 1 times

🗳️ 👤 **Kanasan** 1 year, 4 months ago

**Selected Answer: D**

D. Use a single environment for all containers.

By deploying all microservices in a single environment, you can ensure that they are part of the same virtual network. This makes it easier to manage network policies and allows the services to communicate more easily with each other. Additionally, using a single environment can simplify the configuration for writing logs to the same Log Analytics workspace.

upvoted 2 times

🗳️ 👤 **AashirwadSinha** 1 year, 4 months ago

My vote goes to D

upvoted 1 times

🗳️ 👤 **JeyTienJey** 1 year, 4 months ago

I think option D is correct.

A Container Apps environment is a secure boundary around groups of container apps that share the same virtual network and write logs to the same logging destination.

<https://learn.microsoft.com/en-us/azure/container-apps/environment>



upvoted 2 times

🗳️ 👤 **AlexMMReis** 1 year, 5 months ago

**Selected Answer: D**

seems the only one to make sense D

upvoted 2 times

  **lyggwtwczxnhznebw** 1 year, 5 months ago

**Selected Answer: D**

voting for D as well

upvoted 2 times

HOTSPOT

-

You are developing several microservices to run on Azure Container Apps. External HTTP ingress traffic has been enabled for the microservices.

A deployed microservice must be updated to allow users to test new features. You have the following requirements:

- Enable and maintain a single URL for the updated microservice to provide to test users.
- Update the microservice that corresponds to the current microservice version.

You need to configure Azure Container Apps.

Which features should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area****Requirement**

Single URL for test users

Current microservice activation

Current microservice activation

**Feature**

Revision label  
Revision mode  
Container image  
Container registry

Revision label  
Revision mode  
Container image  
Container registry

**Answer Area****Requirement**

Single URL for test users

Current microservice activation

Current microservice activation

**Feature**

Revision label  
Revision mode  
Container image  
Container registry

Revision label  
Revision mode  
Container image  
Container registry

Suggested Answer:

 **edlima** Highly Voted 1 year, 4 months ago

1. Revision Label

Labels are useful for testing new revisions. For example, when you want to give access to a set of test users, you can give them the label's URL. Then when you want to move your users to a different revision, you can move the label to that revision.

2. Revision Mode:

The revision mode controls whether only a single revision or multiple revisions of your container app can be simultaneously active. You can set your app's revision mode from your container app's Revision management page in the Azure portal, using Azure CLI commands, or in the ARM template.

<https://learn.microsoft.com/en-us/azure/container-apps/revisions>

upvoted 12 times

 **Mpho7** Highly Voted 1 year, 5 months ago

Correct.

Revision labels are most useful when the app is in multiple revision mode.

Revision mode controls whether only a single revision or multiple revisions of your container app can be simultaneously active.

upvoted 6 times

  **bgbgvf** Most Recent 1 year ago

I think given answer is correct.

upvoted 1 times

  **shekhar11** 1 year, 3 months ago

Question in my exam 22sept 2023

upvoted 4 times

  **nikipediaa** 1 year, 3 months ago

me too, today

upvoted 2 times

  **macobuzi** 1 year, 4 months ago

1. Revision Label:

For container apps with external HTTP traffic, labels are a portable means to direct traffic to specific revisions. A label provides a unique URL that you can use to route traffic to the revision that the label is assigned.

2. Revision Mode

The revision mode controls whether only a single revision or multiple revisions of your container app can be simultaneously active.

<https://learn.microsoft.com/en-us/azure/container-apps/revisions>

upvoted 4 times

HOTSPOT

-

You plan to develop an Azure Functions app with an HTTP trigger.

The app must support the following requirements:

- Event-driven scaling
- Ability to use custom Linux images for function execution

You need to identify the app's hosting plan and the maximum amount of time that the app function can take to respond to incoming requests.

Which configuration setting values should you use? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

#### Configuration setting

Hosting plan

#### Value

Consumption  
Dedicated  
Premium

Timeout value

230 seconds  
10 minutes  
unlimited

### Answer Area

#### Configuration setting

Hosting plan

#### Value

Consumption  
Dedicated  
**Premium**

Timeout value

**230 seconds**  
10 minutes  
unlimited

Suggested Answer:

 **Ciupaz**  1 year, 7 months ago

Correct:

The Premium plan supports event-driven scaling and allows to use custom Linux images.

The default timeout for Azure Functions on the Consumption and Premium plans is 5 minutes (300 seconds), and 230 is a good fit.

upvoted 20 times

 **[Removed]** 1 year, 7 months ago

"Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. "

ref: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale>

upvoted 9 times

  **130nk3r5** 1 year, 6 months ago

There is not HTTP trigger in this one ?

upvoted 1 times

  **130nk3r5** 1 year, 6 months ago

Ignore me. I'm stupid

upvoted 6 times

  **1CY1** 1 year ago

but Dedicated does not do both Event Driven scaling and Linux hosting.

upvoted 1 times

  **1CY1** 1 year ago

I thought Dedicated now supports Linux.

upvoted 1 times

  **1CY1** 1 year ago

but Dedicated does not do both Event Driven scaling and Linux hosting.

upvoted 1 times

  **130nk3r5** 1 year, 6 months ago

Got this today.

Went with answer here.

Score 927

upvoted 3 times

  **ciamp**  10 months, 2 weeks ago

Hosting plan:

Premium: The Premium plan supports custom Linux images and event-driven scaling. It provides more flexibility and scalability compared to the Consumption and Dedicated plans.

Timeout value:

Unlimited: The Premium plan allows for unlimited timeout settings, which is beneficial for long-running requests. In contrast, the Consumption plan has a default maximum timeout of 5 minutes (which can be extended to 60 minutes with some configurations) and the Dedicated plan typically has a maximum timeout of 60 minutes.

Correct Configuration Settings:

Hosting plan: Premium

Timeout value: Unlimited

upvoted 6 times

  **Fadoua24**  5 days, 6 hours ago

la respuesta correcta es:

premium (por image linux) y unlimited (Http sin restriccion en premium)

upvoted 1 times

  **J\_obando13** 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

  **Simanta\_M** 8 months, 2 weeks ago

Premium, 230 seconds

upvoted 1 times


  **ciamp** 10 months, 2 weeks ago

1. Premium

2. Unlimited

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale>

upvoted 3 times



  **overhill** 9 months, 1 week ago

this might be incorrect

this affects all tiers



"Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer. For longer processing times, consider using the Durable Functions async pattern or defer the actual work and return an immediate response."

upvoted 3 times

  **prabhjot** 10 months, 3 weeks ago

perfect ans - Premium and @30 - <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale> ( just read about flexi plan also it is getting into GA soon)

upvoted 1 times

  **8ac3742** 10 months, 4 weeks ago

Both Dedicated plan and Premium plan support the custom Linux image for Azure function, however, Dedicated plan doesn't support responsive event-driven autoscaling, it only supports manual scaling and autoscaling slower than autoscaling slower than the premium plan.

upvoted 1 times

  **jobolesoniha1** 1 year, 3 months ago

Premium Plan - ✓ You want to provide a custom Linux image on which to run your functions.

230 Sec - ✓ Regardless of the function app timeout setting, 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer.

upvoted 2 times



  **InversaRadice** 1 year, 4 months ago

Well maximum timeout for premium is 10 minutes, so u re all wrong...

Premium

10 Minutes

upvoted 1 times

  **AlbertoBT** 1 year, 7 months ago

Premium - Unlimited

1 Premium. Is the only one that supports CUSTOM linux images and Event-driven (is not talking about code function runing under linux, in that case could be Consumption)

2. Unlimited - Is talking about set the maximun timeout (has sense form the exam perspective to know witch is the max. and not only a valid one)

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#overview-of-plans>

Consider the Azure Functions Premium plan in the following situations:

+ You want to provide a custom Linux image on which to run your functions.

upvoted 4 times

  **AlbertoBT** 1 year, 7 months ago

UPDATE:

Timeout:

Premium max -> unbounded

Dedicated max -> unbounded

Consumption max -> 10 minutes

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits>

But 230 seconds is the maximum amount of time that an HTTP triggered function can take to respond to a request. This is because of the default idle timeout of Azure Load Balancer.

<https://learn.microsoft.com/en-us/answers/questions/758823/httptrigger-response-timeout>

Premium unlimited could be a valid answer for a Durable Function

upvoted 4 times

  **dddddd111** 1 year, 7 months ago

Correct answers.

upvoted 1 times

  **Schwopak** 1 year, 7 months ago

I think its:

Consumption

10 minutes

They ask about the MAXIMUM timeout (10 minutes for consumption plan).

Consumption plan supports event driven scaling and supports custom Linux images as well.

Maximum Timeout for Premium plan is unbounded according to the tables.

But that would mean, Premium Plan + unlimited would also be a correct answer...

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale>

upvoted 6 times

  **Schwopak** 1 year, 7 months ago

Have to correct my answer: Consumption plan does not allow for CUSTOM Linux images.

So the correct answer would have to be:

Premium

unlimited

again source:

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale>

upvoted 7 times



HOTSPOT

-

You develop a Python application for image rendering. The application uses GPU resources to optimize rendering processes.

You have the following requirements:

- The application must be deployed to a Linux container.
- The container must be stopped when the image rendering is complete.
- The solution must minimize cost.

You need to deploy the application to Azure.

### Answer Area

#### Environment configuration

Compute target

#### Configuration value

Azure Container Instances  
Azure Kubernetes Service  
Azure Container Apps  
Azure App Service

Container termination

Restart policy  
Environment variable  
System-assigned Managed identity  
User-assigned Managed identity

### Answer Area

#### Environment configuration

Compute target

#### Configuration value

Azure Container Instances  
**Azure Kubernetes Service**  
Azure Container Apps  
Azure App Service

Container termination

**Restart policy**  
Environment variable  
System-assigned Managed identity  
User-assigned Managed identity

Suggested Answer:

 **AlbertoBT**  1 year, 1 month ago

ACI - Because the GPU usage. Kubernetes can manage ACIs but is not a 'compute target' and it will increment the cost.

"The container instances in the group can access one or more NVIDIA Tesla GPUs while running container workloads such as CUDA and deep learning applications."

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-gpu>

Restart Policy - To stop the container after the execution (In fact is avoiding to restart it after a succeeded execution)

"Set an appropriate restart policy for the container instance, depending on whether the command-line specifies a long-running task or a run-once task. For example, a restart policy of Never or OnFailure is recommended for a run-once task."

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-start-command#command-line-guidelines>

upvoted 14 times

👤 **onlyforheros** Highly Voted 9 months, 2 weeks ago

Got it in exam on 13.03.2024. Score: 910.

Answer ACI, restart policy

upvoted 10 times

👤 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 2 times

👤 **jobolesoniha1** 9 months, 1 week ago

Azure Container Instance - Since scaling is not an requirement, the right choice would be to use ACI instead of ACA or Kubernetes.

upvoted 2 times

👤 **manopeydakon** 11 months, 3 weeks ago

To deploy a Python application for image rendering with the specified requirements on Azure, you can use Azure Container Instances (ACI). ACI provides an easy and cost-effective way to run containers in the cloud. Here are the options you can choose:

Compute Value:

Choose the ACI option for running your Linux container.

Container Termination:

Since you want the container to be stopped when the image rendering is complete to minimize cost, you can use the "Single" container group mode in Azure Container Instances. This mode is suitable for scenarios where you want to run a single container task and stop it when it's done.

This command creates a single container instance (--restart-policy Never) and stops it when the image rendering is complete, minimizing cost. Adjust the resource limits (--cpu and --memory) based on your application's requirements.

upvoted 4 times

👤 **azrlm** 1 year ago

Will this question be on the exam?

upvoted 1 times

👤 **SachinV** 1 year, 1 month ago

As suggested ACI is preferred for one time or demand script and jobs execution whereas as ACA is preferred for deploying container continually running full lifecycle apps, I am leaning towards ACI.

upvoted 2 times

👤 **dddddd111** 1 year, 1 month ago

If I choose C# in the exam, will this question still appear?

upvoted 1 times

👤 **junkz** 1 year, 1 month ago

given that AKS is not in the exam curriculum , only ACA and ACI are left. out of which, only ACI supports GPU (though still in preview). As for the container termination-> restart policy set to never should be true

upvoted 6 times

👤 **Schwopak** 1 year, 1 month ago

I agree with ACI

Since its a new question the preview thing makes sense as well.

source:

<https://learn.microsoft.com/en-gb/azure/container-instances/container-instances-overview>

upvoted 1 times

## HOTSPOT

-

You plan to develop an Azure Functions app with an Azure Blob Storage trigger. The app will be used infrequently, with a limited duration of individual executions.

The app must meet the following requirements:

- Event-driven scaling
- Support for deployment slots
- Minimize costs

You need to identify the hosting plan and the maximum duration when executing the app.

Which configuration setting values should you use? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Configuration setting	Value
Hosting plan	<div>Consumption</div> <div>Dedicated</div> <div>Premium</div>
Maximum execution time	<div>230 seconds</div> <div>10 minutes</div> <div>unlimited</div>

**Answer Area**

Suggested Answer:

Configuration setting	Value
Hosting plan	<div>Consumption</div> <div>Dedicated</div> <div>Premium</div>
Maximum execution time	<div>230 seconds</div> <div>10 minutes</div> <div>unlimited</div>

🗨️ 👤 Ciupaz Highly Voted 1 year, 7 months ago

Correct.

"used infrequently" -> Consumption

And 10mins is a good choice.

upvoted 13 times

🗨️ 👤 1CY1 1 year ago

Consumption supports :



- 2 deployment slots.
- minimizes costs for infrequent use.
- event driven scaling.

Timeout:

Consumption max : 10 minutes

Answer => Consumption, 10 minutes.

upvoted 2 times

  **1CY1** 11 months, 1 week ago

Sorry that should be execution time not timeout.

upvoted 1 times


  **Schwopak** 1 year, 7 months ago

Agree though my reasoning for Consumption stems from "minimize cost".

Max timeout (computation time) for consumption plan azure functions is 10 min.

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale>

upvoted 2 times

  **AlbertoBT**  1 year, 7 months ago

Right Answer: Premium - 230

Event-Driven requirement:

Is only supported by Consumption and Premium

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#scale>

Slot requirement:

Is only supported by Dedicated and Premium

"Function apps running under the Apps Service plan may have multiple slots, while under the Consumption plan only one slot is allowed."

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-deployment-slots>

Minimize costs:

Is not relevant because there is just one service plan that meets the previous requirements

upvoted 11 times

  **1CY1** 1 year ago

2 deployment slots are supported on a consumption plan.


<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits>

upvoted 1 times

  **raymond\_abcd** 1 year, 4 months ago

Function apps running in a Consumption plan have a single extra slot for staging.

upvoted 2 times

  **SSR999** 1 year, 5 months ago

Consumption plan also supports 2 slots..<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits>

Answer; Consumption - 10 min

upvoted 8 times

  **AlbertoBT** 1 year, 7 months ago

Timeout:

Premium and Dedicated max -> unbounded

Consumption max -> 10 minutes

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits>

230 seconds is the maximum amount of time that an HTTP triggered function

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#timeout>

Premium unlimited could be a good choice for a Durable Function but is for "a limited duration of individual executions" so the minimum value can fit.

upvoted 2 times

  **130nk3r5** 1 year, 6 months ago

No HTTP one here...?

I'll go with Premium and unlimited here. Minimizing costs here is the issue ?!?

This is a stupid question.

upvoted 1 times

🗨️ 👤 **f09229d** Most Recent 3 months ago

Gone in 3/2025 exam with proposed answer 907/1000 (Consumption and 10 min)

upvoted 3 times

🗨️ 👤 **dino\_co** 1 year, 3 months ago

Hosting Plan: Consumption - The Consumption plan automatically scales based on demand, supports deployment slots, and you only pay for compute resources when your functions are running. This makes it a cost-effective choice for applications that are used infrequently.

Maximum Execution Time: 10 minutes. In the Consumption plan, function executions time out after five minutes by default, but this can be set to a maximum of ten minutes.

upvoted 2 times

🗨️ 👤 **FeriAZ** 1 year, 4 months ago

Hosting Plan: Consumption

Maximum Execution Time: 10 minutes

upvoted 3 times

🗨️ 👤 **raymond\_abcd** 1 year, 4 months ago

Consumption plan has 2 slots, see: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale#service-limits>

Max Timeout is 10 minutes (when HTTP triggered 230 secs)

upvoted 1 times

🗨️ 👤 **kumarts** 1 year, 5 months ago

Consumption - 10 minutes is right...Event driven is supported in Consumption and Premium.

Consumption plan also supports 2 slots. Refer weblink - <https://learn.microsoft.com/en-us/azure/azure-functions/functions-scale>

upvoted 1 times

You are developing an ASP.NET Core app hosted in Azure App Service.

The app requires custom claims to be returned from Microsoft Entra ID for user authorization. The claims must be removed when the app registration is removed.

You need to include the custom claims in the user access token.

What should you do?

- A. Require the `https://graph.microsoft.com/.default` scope during authentication.
- B. Configure the app to use the OAuth 2.0 authorization code flow.
- C. Implement custom middleware to retrieve role information from Azure AD.
- D. Add the groups to the `groupMembershipClaims` attribute in the app manifest.
- E. Add the roles to the `appRoles` attribute in the app manifest.

**Suggested Answer: B**

Community vote distribution

E (100%)

🗳️ 👤 **kygukyzo** 6 months, 3 weeks ago

**Selected Answer: E**

Correct

upvoted 1 times

🗳️ 👤 **FeriAZ** 1 year, 4 months ago

**Selected Answer: E**

Azure Active Directory (Azure AD) supports adding custom roles to an application's manifest, which can then be assigned to users or groups. When a user is authenticated, these roles are included in the token as claims. This approach allows for fine-grained access control within your application based on these role assignments. Moreover, when the application registration is deleted, these roles and corresponding claims automatically cease to exist, fulfilling the requirement that the claims must be removed when the app registration is removed.

upvoted 4 times

🗳️ 👤 **AzDeveloper** 1 year, 5 months ago

**Selected Answer: E**

E not D because of this condition

"The claims must be removed when the app registration is removed."

upvoted 3 times

🗳️ 👤 **AzDeveloper** 1 year, 5 months ago

E not D because of this condition

"The claims must be removed when the app registration is removed."

upvoted 1 times

🗳️ 👤 **manopeydakon** 1 year, 5 months ago

To include custom claims in the user access token from Microsoft Identity for user authorization, you should:

E. Add the roles to the `appRoles` attribute in the app manifest.

Explanation:

In the Azure AD app manifest, you can define custom roles using the `appRoles` attribute.

These roles can then be assigned to users, and the associated claims will be included in the user's token.

Ensure that the `appRoles` attribute in the app manifest includes the necessary roles with associated claims, and assign these roles to users accordingly.

This approach allows you to customize the claims included in the user's access token when they authenticate with Microsoft Identity.

upvoted 1 times

🗳️ 👤 **Swekker** 1 year, 5 months ago

Selected Answer: E

AppRoles is the way to go.

<https://learn.microsoft.com/en-us/entra/identity-platform/howto-add-app-roles-in-apps#app-roles-vs-groups>

upvoted 3 times

  **1CY1** 1 year ago

Setting the Microsoft Entra app manifest appRoles attribute.

<https://learn.microsoft.com/en-us/entra/identity-platform/reference-app-manifest#approles-attribute>

upvoted 1 times

You are developing a microservice to run on Azure Container Apps for a company. External HTTP ingress traffic has been enabled.

The company requires that updates to the microservice must not cause downtime.


You need to deploy an update to the microservices.

What should you do?

- A. Enable single revision mode.
- B. Use multiple environments for each container.
- C. Use a private container registry and single image for all containers.
- D. Use a single environment for all containers.
- E. Enable multiple revision mode.

**Suggested Answer: A**

Community vote distribution



🗳️ 👤 **AzDeveloper** Highly Voted 🏆 1 year, 5 months ago

**Selected Answer: A**

Zero downtime deployment using single mode

<https://learn.microsoft.com/en-us/azure/container-apps/revisions#zero-downtime-deployment>

upvoted 12 times

🗳️ 👤 **1CY1** 1 year ago

I agree with Kvadraciuks.

See : <https://learn.microsoft.com/en-us/azure/container-apps/revisions#work-with-multiple-revisions>

The question talks about development which to me implies that you are going to want to stage deployment of new revisions and be able to monitor traffic splitting using multiple revisions.

Answer => E. Enable multiple revision mode.

upvoted 1 times

🗳️ 👤 **J\_obando13** Most Recent 🔍 1 month, 2 weeks ago

**Selected Answer: A**

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

🗳️ 👤 **Imobiletun** 1 month, 4 weeks ago

**Selected Answer: E**

E. Enable multiple revision mode.

upvoted 1 times

🗳️ 👤 **Iaminall** 2 months ago

**Selected Answer: E**

E is the correct one imho

upvoted 1 times

🗳️ 👤 **Iaminall** 3 months ago

**Selected Answer: E**

Enable multiple revision mode.

upvoted 2 times

🗳️ 👤 **Iaminall** 4 months ago

**Selected Answer: E**

E is the correct one.



upvoted 1 times

🗨️ 👤 **5691024** 5 months ago

**Selected Answer: E**

My guess is E?

The docs are very clear that both Single- and Multi-Revision are zero-downtime:

<https://learn.microsoft.com/en-us/azure/container-apps/revisions#zero-downtime-deployment>

...so there's no clear answer just from that.

My thinking is: Even if the container ready-checks pass, the deployed code itself could still be flawed and result in a broken application. Being able to validate the new version using multi-revision mode is a safer option than just sending it. Is a broken app technically 'downtime'? Debatable. Should a broken app be caught before the prod deploy? Certainly. But I don't see any other way to favour one option over the other.

Of course, is that what Mr. Microsoft wants...? I have no idea.

upvoted 4 times

🗨️ 👤 **VithobaS** 5 months, 1 week ago

**Selected Answer: E**

Multiple Revision Mode:

- In multiple revision mode, Azure Container Apps creates a new revision each time a deployment is updated.
- Traffic can be split between the current and new revisions during the update process, ensuring zero downtime.
- Once the new revision is stable, traffic can be shifted fully to the new revision, completing the update.

upvoted 1 times

🗨️ 👤 **NewChampion** 7 months ago

**Selected Answer: A**

Single Revision Mode has zero downtime.

upvoted 1 times

🗨️ 👤 **Dkijc** 8 months, 1 week ago

A

Source: <https://learn.microsoft.com/en-us/azure/container-apps/revisions#zero-downtime-deployment>

upvoted 1 times

🗨️ 👤 **ciamp** 10 months, 2 weeks ago

E

Single Revision Mode:

In single revision mode, only one revision of the application is active at any given time. When you deploy a new revision, it replaces the current active revision.

This means that during the deployment process, the old revision is taken offline while the new one is brought up. This can cause a brief period where the application might be unavailable or not fully functional, resulting in potential downtime.

Multiple Revision Mode:

This mode allows multiple revisions to run concurrently. When deploying a new revision, it can coexist with the old revision. Traffic can be routed gradually from the old revision to the new one, ensuring a smooth transition without downtime.

upvoted 1 times

🗨️ 👤 **Kvadraciuks** 1 year, 2 months ago

**Selected Answer: E**

E. Enable multiple revision mode:

Multiple revision mode allows for more than one version of the container app to be active at the same time. This is essential for zero-downtime deployments because you can deploy a new revision without taking down the old one, then gradually shift traffic to the new revision using traffic splitting.

upvoted 4 times

🗨️ 👤 **azpardha** 1 year, 3 months ago

**Selected Answer: A**

single revision mode

upvoted 1 times

🗨️ 👤 **Sachin39** 1 year, 5 months ago

**Selected Answer: A**

In single revision mode, Container Apps ensures your app doesn't experience downtime when creating a new revision. The existing active revision isn't deactivated until the new revision is ready.

If ingress is enabled, the existing revision continues to receive 100% of the traffic until the new revision is ready.

<https://learn.microsoft.com/en-us/azure/container-apps/revisions>

upvoted 4 times

🗨️ 👤 **manopeydaikon** 1 year, 5 months ago

E. Enable multiple revision mode.

Explanation:

Azure Container Apps supports multiple revision mode, allowing you to deploy updates without causing downtime.

With multiple revision mode, you can deploy new revisions of your microservice while keeping the previous revision running.

This approach enables a smooth transition without interruption to your application.

So, choose option E to enable multiple revision mode when deploying updates to your microservice on Azure Container Apps.

upvoted 2 times

🗨️ 👤 **Woksi** 1 year, 4 months ago

The task is to deploy "an update". Maybe they are trying to signal that there is no need to allow the old version to continue running when the update is completed. A/B testing is not mentioned.

I agree that it would solve the ask - but Single Revision Mode does so as well, without the need for manual setup of Traffic Splitting, and without having to manually clean up old revisions.

upvoted 1 times

🗨️ 👤 **AlpeshD** 1 year, 5 months ago

**Selected Answer: A**

[https://www.google.co.nz/search?q=which+revision+mode+in+azure+for+zero+downtime&safe=active&sca\\_esv=596418299&ei=yyObZa-sM9GYseMP2IOmqA0&ved=0ahUKEwjvvdT3p8yDaxVRTGwGHdiBCdUQ4dUDCBA&uact=5&oq=which+revision+mode+in+azure+for+zero+downtime&gs\\_lp=EgAEBwgIKEAAYRxiWBBiWA8ICBRAhGKABwgIEECEYFeIDBBgAIEGIBgGQBgg&sclient=gws-wiz-serp](https://www.google.co.nz/search?q=which+revision+mode+in+azure+for+zero+downtime&safe=active&sca_esv=596418299&ei=yyObZa-sM9GYseMP2IOmqA0&ved=0ahUKEwjvvdT3p8yDaxVRTGwGHdiBCdUQ4dUDCBA&uact=5&oq=which+revision+mode+in+azure+for+zero+downtime&gs_lp=EgAEBwgIKEAAYRxiWBBiWA8ICBRAhGKABwgIEECEYFeIDBBgAIEGIBgGQBgg&sclient=gws-wiz-serp)

upvoted 1 times

🗨️ 👤 **Circle089** 1 year, 5 months ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/azure/container-apps/revisions#zero-downtime-deployment>

upvoted 2 times

HOTSPOT

-

A company uses Azure Container Apps. A container app named App1 resides in a resource group named RG1.

The company requires testing of updates to App1.

You enable multiple revision modes on App1.

You need to ensure traffic is routed to each revision of App1.

How should you complete the code segment?

NOTE: Each correct selection is worth one point.

### Answer Area

```
az 

container



containerapp



network



resource



app



connection



ingress



revision

 traffic set \  
  
--name App1 \  
--resource-group RG1 \  
--revision-weight <REVISION_1>=80 <REVISION_2>=20
```

### Answer Area

Suggested Answer:

```
az 

container



containerapp



network



resource



app



connection





ingress



revision

 traffic set \  
  
--name App1 \  
--resource-group RG1 \  
--revision-weight <REVISION_1>=80 <REVISION_2>=20
```

 **1989jay**  1 year, 5 months ago

Correct. containerapp, ingress

Source: <https://learn.microsoft.com/en-us/azure/container-apps/traffic-splitting?pivots=azure-cli>

upvoted 15 times

 **1CY1** 1 year ago

Configure traffic splitting between revisions using the az containerapp ingress traffic set command.

```
az containerapp ingress traffic set --name <APP_NAME>  
--resource-group <RESOURCE_GROUP>  
--revision-weight <REVISION_1>=50 <REVISION_2>=50
```

upvoted 2 times

  **azpardha** Most Recent 1 year, 3 months ago

Correct. containerapp, ingress

upvoted 1 times

  **kysmf0100101** 1 year, 5 months ago

CORRECT

upvoted 2 times

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 Container Apps app and disable ingress on the container app.

Users report that they are unable to access the container app. You investigate and observe that the app has scaled to 0 instances.

You need to resolve the issue with the container app.

Solution: Enable ingress, create an HTTP scale rule, and apply the rule to the container app.


Does the solution meet the goal?

- A. Yes
- B. No

**Suggested Answer: A**

Community vote distribution

B (100%)

 **MarcosAn** Highly Voted 1 year, 5 months ago

**Selected Answer: B**

Must be carefull with Default Scale Rule: <https://learn.microsoft.com/en-us/azure/container-apps/scale-app?pivots=azure-cli#default-scale-rule>  
Make sure you create a scale rule or set minReplicas to 1 or more if you don't enable ingress. If ingress is disabled and you don't define a minReplicas or a custom scale rule, then your container app will scale to zero and have no way of starting back up.

upvoted 9 times


 **vizay** Most Recent 2 months ago

**Selected Answer: B**

No, the solution does not meet the goal.

Here's why: Enabling ingress and setting up a scale rule is a good idea, but you also need to make sure the app is set up to scale automatically based on traffic. Just enabling ingress won't fix the issue of the app scaling to 0 instances.

upvoted 1 times

 **8ac3742** 10 months, 4 weeks ago

**Selected Answer: A**

enabled ingress->create an HTTP scale rule to extend instances based on incoming traffic

upvoted 3 times

 **Chriscrea8** 1 year, 3 months ago

No, the solution does not meet the goal.

Enabling ingress and creating an HTTP scale rule won't directly address the issue of the container app scaling down to 0 instances. Enabling ingress only allows external traffic to reach the container app, but it doesn't guarantee that the app will scale up or have instances available to handle incoming requests.

To resolve the issue of the container app scaling down to 0 instances, you need to adjust the scaling settings to ensure that there are always instances available to handle incoming requests, even when traffic is low. This may involve modifying the minimum instance count or configuring auto-scaling rules based on metrics such as CPU usage or incoming requests.

So, the correct answer is B. No.

upvoted 2 times

🗨️ 👤 **Mahesh1222** 1 year, 3 months ago

Though Enable Ingress, Set Scale rule to Http is correct, it will work based on Http traffic coming in when the Instance is already running. But it doesn't address the problem mentioned in the question - where it already scaled down to 0 instances. So to address this first, we need to ensure that there is at least one instance of the container app running to handle incoming requests. This may involve adjusting the scaling settings, configuring minimum instance limits, or implementing a health check mechanism to keep at least one instance running at all times, even during periods of low traffic.

upvoted 3 times

🗨️ 👤 **jobolesonihal** 1 year, 3 months ago

**Selected Answer: B**

The HTTP scale rule in Azure App Service allows you to automatically scale your app based on HTTP traffic. However, if the Azure container app is scaled down to 0 instances, the HTTP scale rule will not work because there are no instances to scale.

In other words, the HTTP scale rule only applies when there are active instances of the app to handle incoming HTTP requests. When the app is scaled down to 0 instances, no requests are being processed, and the scale rule does not have any effect.

Therefore, if you want the HTTP scale rule to work, you need to ensure that there are at least one or more instances of the app running.

upvoted 1 times

🗨️ 👤 **Ushakanth** 1 year, 4 months ago

**Selected Answer: B**

Though Enable Ingress, Set Scale rule to Http is correct, it will work based on Http traffic coming in when the Instance is already running. But it doesn't address the problem mentioned in the question - where it already scaled down to 0 instances. So to address this first, we need to ensure that there is at least one instance of the container app running to handle incoming requests. This may involve adjusting the scaling settings, configuring minimum instance limits, or implementing a health check mechanism to keep at least one instance running at all times, even during periods of low traffic.

upvoted 2 times

🗨️ 👤 **emysa341** 1 year, 5 months ago

A is correct since scale rule can be HTTP

<https://learn.microsoft.com/en-us/azure/container-apps/scale-app?pivots=azure-cli#:~:text=Scale%20rules%20are%20implemented%20as%20HTTP%2C%20TCP%2C%20or%20custom.>

upvoted 1 times

🗨️ 👤 **FeriAZ** 1 year, 5 months ago

Answer is A:

Enabling Ingress: This would allow external traffic to reach the container app, which is necessary for users to access the app from outside the Azure network.

Creating an HTTP Scale Rule: Azure Container Apps can automatically scale based on various metrics, including HTTP traffic. An HTTP scale rule would allow the app to scale up based on the incoming HTTP traffic.

Applying the Rule to the Container App: Implementing the scale rule ensures that the container app scales appropriately in response to traffic, preventing it from scaling down to 0 instances when there is demand.

upvoted 3 times

🗨️ 👤 **ryuhei** 1 year, 5 months ago

**Selected Answer: B**

Answer is B

upvoted 2 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

No!

Enabling ingress and creating an HTTP scale rule won't directly address the issue of the app having scaled to 0 instances. Ingress relates to the network traffic coming into the app, and scale rules define when and how the app scales based on certain conditions. However, they won't directly bring the app back to running instances if it has scaled down to 0.

To address the issue of the app having 0 instances, you might need to investigate why the app scaled down and take appropriate action, such as adjusting the scaling settings, reviewing resource constraints, or identifying any issues in the app itself.

upvoted 2 times

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 Container Apps app and disable ingress on the container app.

Users report that they are unable to access the container app. You investigate and observe that the app has scaled to 0 instances.

You need to resolve the issue with the container app.

Solution: Enable ingress, create a custom scale rule, and apply the rule to the container app.


Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Community vote distribution



🗳️ **Munwalinwali** 12 months ago

**Selected Answer: B**

#62 is the correct answer

upvoted 1 times

🗳️ **JoaoPelisson** 1 year, 1 month ago

**Selected Answer: B**

Não atende

upvoted 2 times

🗳️ **jobolesonihal** 1 year, 3 months ago

**Selected Answer: B**

Documentation says "Make sure you create a scale rule or set minReplicas to 1 or more if you don't enable ingress. If ingress is disabled and you don't define a minReplicas or a custom scale rule, then your container app will scale to zero and have no way of starting back up."

upvoted 4 times

🗳️ **onlyforheros** 1 year, 4 months ago

**Selected Answer: B**

Creating a custom scale rule won't directly resolve the issue of the container app being inaccessible when scaled to 0 instances.

Only if you set the scale to 1

upvoted 1 times

🗳️ **4d716d6** 1 year, 4 months ago

**Selected Answer: B**

Same as #60

upvoted 1 times

🗳️ **vkalynenko** 1 year, 2 months ago

#60 mentioned Http scale rule, not custom

upvoted 3 times

🗳️ **emysa341** 1 year, 5 months ago

**Selected Answer: A**

A is correct since scale rule can be custom

<https://learn.microsoft.com/en-us/azure/container-apps/scale-app?pivots=azure->

[cli#:~:text=Scale%20rules%20are%20implemented%20as%20HTTP%2C%20TCP%2C%20or%20custom.](#)



upvoted 3 times

  **de8a1d8** 1 year, 4 months ago

Right, page says : Important

Make sure you create a scale rule or set minReplicas to 1 or more if you don't enable ingress. If ingress is disabled and you don't define a minReplicas "or a custom scale rule", then your container app will scale to zero and have no way of starting back up.

upvoted 2 times

  **ryuhei** 1 year, 5 months ago

The problem is the same as #60.

upvoted 3 times

  **MarcosAn** 1 year, 5 months ago

Yes, same cenario as previus question.

upvoted 1 times



  **manopeydaon** 1 year, 5 months ago

No,

Enabling ingress and creating a custom scale rule won't directly address the issue of the app having scaled to 0 instances. Ingress and scale rules are related to managing the app's scalability and handling incoming traffic conditions, but they won't bring the app back to running instances if it has scaled down to 0.

To resolve the issue of the app having 0 instances, you need to investigate why the app scaled down and take appropriate action. Reviewing scaling settings, resource constraints, and identifying any issues in the app itself would be more relevant to resolving this specific problem.

upvoted 4 times

  **Ciupaz** 1 year, 5 months ago

**Selected Answer: A**

Yes, this could be a good solution.

upvoted 1 times



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 Container Apps app and disable ingress on the container app.

Users report that they are unable to access the container app. You investigate and observe that the app has scaled to 0 instances.

You need to resolve the issue with the container app.

Solution: Enable ingress and configure the minimum replicas to 1 for the container app.


Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Community vote distribution



**Mahesh1222** 9 months, 1 week ago

**Selected Answer: A**

The minimum replicas setting=1 ensures that the app always has at least one instance running.  
upvoted 3 times

**sky\_may** 11 months, 1 week ago

Answer is A: From the documentation:

Make sure you create a scale rule or set minReplicas to 1 or more if you don't enable ingress. If ingress is disabled and you don't define a minReplicas or a custom scale rule, then your container app will scale to zero and have no way of starting back up.

<https://learn.microsoft.com/en-us/azure/container-apps/scale-app?pivots=azure-cli#default-scale-rule>  
upvoted 3 times

**sky\_may** 11 months, 1 week ago

From doc: If you want to ensure that an instance of your revision is always running, set the minimum number of replicas to 1 or higher.  
upvoted 1 times

**emysa341** 11 months, 2 weeks ago

**Selected Answer: B**

The correct answer is B since the minimum replica of the scale can be 0

<https://learn.microsoft.com/en-us/azure/container-apps/scale-app?pivots=azure-cli#:~:text=Minimum%20number%20of%20replicas%20per,300>  
upvoted 1 times

**FeriAZ** 11 months, 2 weeks ago

answer is A:

Enable Ingress: Ingress is crucial for allowing external access to your container app. Disabling ingress had led to the app being inaccessible to users. Enabling it will reopen external traffic routes, allowing users to access the app.

Configure the Minimum Replicas to 1: This setting ensures that there is always at least one instance of the app running. By setting the minimum number of replicas to 1, you prevent the app from scaling down to 0 instances, which addresses the issue observed where the app had no running instances.

upvoted 2 times

🗨️ 👤 **ryuhei** 11 months, 3 weeks ago

**Selected Answer: A**

Answer is A Only

upvoted 2 times

🗨️ 👤 **manopeydakon** 11 months, 3 weeks ago

A,

Yes, the solution meets the goal. By enabling ingress and configuring the minimum replicas to 1, you ensure that there is at least one running instance of the container app. This allows users to access the app, as there is at least one instance available to handle incoming requests.

upvoted 2 times

🗨️ 👤 **Ciupaz** 11 months, 4 weeks ago

**Selected Answer: A**

For me this could be a good solution. The minimum replicas setting=1 ensures that the app always has at least one instance running.

upvoted 2 times

## HOTSPOT

-

## Case study

-

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

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

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

## To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

## Background

-

Munson's Pickles and Preserves Farm is an agricultural cooperative corporation based in Washington, US, with farms located across the United States. The company supports agricultural production resources by distributing seeds fertilizers, chemicals, fuel, and farm machinery to the farms.

## Current Environment

-

The company is migrating all applications from an on-premises datacenter to Microsoft Azure. Applications support distributors, farmers, and internal company staff.

## Corporate website

-

- The company hosts a public website located at <http://www.munsonspicklesandpreservesfarm.com>. The site supports farmers and distributors who request agricultural production resources.

## Farms

-

- The company created a new customer tenant in the Microsoft Entra admin center to support authentication and authorization for applications.

## Distributors

-

- Distributors integrate their applications with data that is accessible by using APIs hosted at

<http://www.munsonspicklesandpreservesfarm.com/api> to receive and update resource data.

## Requirements

-

The application components must meet the following requirements:

### Corporate website

-

- The site must be migrated to Azure App Service.
- Costs must be minimized when hosting in Azure.
- Applications must automatically scale independent of the compute resources.
- All code changes must be validated by internal staff before release to production.
- File transfer speeds must improve, and webpage-load performance must increase.
- All site settings must be centrally stored, secured without using secrets, and encrypted at rest and in transit.
- A queue-based load leveling pattern must be implemented by using Azure Service Bus queues to support high volumes of website agricultural production resource requests.

### Farms

-

- Farmers must authenticate to applications by using Microsoft Entra ID.

### Distributors

-

- The company must track a custom telemetry value with each API call and monitor performance of all APIs.
- API telemetry values must be charted to evaluate variations and trends for resource data.

### Internal staff

-

- App and API updates must be validated before release to production.
- Staff must be able to select a link to direct them back to the production app when validating an app or API update.
- Staff profile photos and email must be displayed on the website once they authenticate to applications by using their Microsoft Entra ID.

### Security

-

- All web communications must be secured by using TLS/HTTPS.
- Web content must be restricted by country/region to support corporate compliance standards.
- The principle of least privilege must be applied when providing any user rights or process access rights.
- Managed identities for Azure resources must be used to authenticate services that support Microsoft Entra ID authentication.

### Issues

-

### Corporate website

-

- Farmers report HTTP 503 errors at the same time as internal staff report that CPU and memory usage are high.
- Distributors report HTTP 502 errors at the same time as internal staff report that average response times and networking traffic are high.
- Internal staff report webpage load sizes are large and take a long time to load.
- Developers receive authentication errors to Service Bus when they debug locally.

Distributors

- Many API telemetry values are sent in a short period of time. Telemetry traffic, data costs, and storage costs must be reduced while preserving a statistically correct analysis of the data points sent by the APIs.

You need to configure App Service to support the corporate website migration.

Which configuration should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Azure App Service configuration

#### Configuration setting

App Service plan

#### Configuration value

Basic  
Standard  
Premium  
Isolated

Code change validation feature

Deployment slot  
Custom container  
Domain certificate  
Deployment credentials

### Azure App Service configuration

#### Configuration setting

App Service plan


#### Configuration value

Basic  
Standard  
Premium  
Isolated

Code change validation feature

Deployment slot  
Custom container  
Domain certificate  
Deployment credentials

Suggested Answer:

 **priomkhan** Highly Voted 1 year, 5 months ago

Basic App Service Plan does not have multiple Deployment slot. Answer should be Standard and Deployment Slot.

upvoted 27 times

 **1CY1** 1 year ago


Answer => Standard.

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits#app-service-limits>



Basic does not have auto scaling or deployment slots.

Standard has auto scaling and deployment slots.

upvoted 2 times

 **sky\_may** 1 year, 5 months ago

correct: <https://learn.microsoft.com/en-us/azure/app-service/deploy-staging-slots?source=recommendations&tabs=portal>  
upvoted 1 times

  **harbox** 1 year, 3 months ago

Where you indicate it clearly states:

"...you can use a separate deployment slot instead of the default production slot when you're running in the Standard, Premium, or Isolated App Service plan tier."

upvoted 1 times

  **InversaRadice** 1 year, 4 months ago

Absolutely wrong, because "applications must autoscale", basic plan doesn't allowe this, the answer is Standard:

<https://learn.microsoft.com/en-us/azure/app-service/manage-automatic-scaling?tabs=azure-portal>

upvoted 2 times

  **manopeydakon**  1 year, 5 months ago

Standard and Deployment Slot.

upvoted 12 times

  **J\_obando13**  1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)



upvoted 1 times

  **Ira89** 3 months, 2 weeks ago

Standard

Deployment Slot

upvoted 1 times

  **FeriAZ** 1 year, 4 months ago

App Service Plan: Standard. It meets the need for auto-scaling and cost efficiency while providing the necessary features such as SSL and custom domains.

Code Change Validation Feature: Deployment Slot. Enables the company to validate all code changes in a pre-production environment, ensuring stability and performance before changes are made live.

upvoted 6 times

## DRAG DROP

-

You have an Azure Cosmos DB for NoSQL API account named account1 and a database named db1. An application named app1 will access db1 to perform read and write operations.

You plan to modify the consistency levels for read and write operations performed by app1 on db1.

You must enforce the consistency level on a per-operation basis whenever possible.

You need to configure the consistency level for read and write operations.

Which locations should you configure? To answer, move the appropriate locations to the correct operations. You may use each location once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

## Locations to configure operations

db1

app1

account1

## Locations to configure operations

Operation	Location
Read	
Write	

**Suggested Answer:**

Operation	Location
Read	app1
Write	account1

**de55d65** Highly Voted 7 months, 1 week ago

Hency only read can be overwritten by the SDK I'd go with

read: app1

write: account1

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level>

upvoted 5 times

**Piddi** 6 months ago

Consistency can only be relaxed at the SDK instance or request level. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account.

upvoted 1 times

**Piddi** 6 months ago

Overriding the default consistency level only applies to reads within the SDK client. An account configured for strong consistency by default will still write and replicate data synchronously to every region in the account. When the SDK client instance or request overrides this with Session or weaker consistency, reads will be performed using a single replica.

upvoted 1 times

**Alekeyh** Most Recent 3 months ago

Answer is ok.

Reads are flexible, writes are firm, only can be changed in the account.

upvoted 1 times

🗨️ 👤 **de55d65** 7 months, 3 weeks ago

Shouldn't write be on the account hence the write consistency level cannot be set within the app?

upvoted 1 times

🗨️ 👤 **Mattt** 7 months, 3 weeks ago

--per operation--

It should be app1 for read and write

upvoted 2 times

🗨️ 👤 **alvingeo** 8 months ago

It is app1 for both Read and Write operations, as per-operation consistency can be set from the client application level (app1) in this scenario.

upvoted 4 times

🗨️ 👤 **fc61c73** 8 months, 1 week ago

it is app1 for both, it says per operation that mean from client location in this case app1.

upvoted 3 times



HOTSPOT

-

You are creating an Azure Functions app project in your local development environment by using Azure Functions Core Tools.

You must create the project in either Python or C# without using a template.

You need to specify the command and its parameter required to create the Azure Functions app project.

Which command and parameter should you specify? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

## Core Tools syntax elements

### Syntax element Value

Command

	▼
func new func imt func azure	

Parameter

	▼
--language --worker-runtime --target-framework	

Suggested Answer:

### Core Tools syntax elements


#### Syntax element Value

Command

	▼
func new func imt func azure	

Parameter

	▼
--language --worker-runtime --target-framework	

 **NikolaMicic** Highly Voted 8 months, 2 weeks ago

"func imt" must be "func init"

if you need a new project without a template in python or c#, the correct answer:

func init, --worker-runtime

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-core-tools-reference>

upvoted 19 times


 **PereCastor37** 6 months, 2 weeks ago

agreed, because "func new" is to create a new function based on a template, and the sentence said "without using a template"  
<https://learn.microsoft.com/en-us/azure/azure-functions/functions-core-tools-reference?tabs=v2#func-new>  
upvoted 1 times

  **amircoka** Most Recent 7 months, 2 weeks ago

- func new - creates new function within current project. But we need to create new project, so it should be:

- func init  
- --worker-runtime  
upvoted 1 times

  **Mattt** 8 months, 1 week ago

func init <projectName> --python



func init <projectName> --cs  
upvoted 1 times

  **Mattt** 7 months, 3 weeks ago

ignore my answer. it should be :

func init  
--worker-runtime

--worker-runtime Sets the language runtime for the project. Supported values are: csharp, dotnet, dotnet-isolated, javascript,node (JavaScript), powershell, python, and typescript. For Java, use Maven. To generate a language-agnostic project with just the project files, use custom. When not set, you're prompted to choose your runtime during initialization.  
upvoted 1 times

  **NikolaMicic** 8 months, 2 weeks ago

"func imt" must be "func init"

if you need a new project without a template in python or c#, the correct answer:

func init, --worker-runtime  
upvoted 4 times

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 App Service plan named APSPPlan1 set to the Basic B1 pricing tier. APSPPlan1 contains an App Service web app named WebApp1.

You plan to enable schedule-based autoscaling for APSPPlan1.

You need to minimize the cost of running WebApp1.

Solution: Scale down APSPPlan1 to the Shared pricing tier.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Community vote distribution

B (100%)

  **examtopicsLogin123** 6 months, 3 weeks ago

**Selected Answer: B**



You need the Autoscale to setup the Schedule-based scaling

<https://learn.microsoft.com/en-us/azure/app-service/manage-automatic-scaling>

and the Autoscale is available from Standard and above

<https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits#app-service-limits>

upvoted 4 times

  **examtopicsLogin123** 6 months, 3 weeks ago

One more confirmation: "Not all pricing tiers support autoscaling. The development pricing tiers are either limited to a single instance (the F1 and D1 tiers), or they only provide manual scaling (the B1 tier). If you've selected one of these tiers, you must first scale up to the S1 or any of the P level production tiers."

<https://learn.microsoft.com/en-us/training/modules/scale-apps-app-service/4-autoscale-app-service>

upvoted 1 times

  **Mattt** 7 months, 3 weeks ago

**Selected Answer: B**

No

The Shared pricing tier does not support autoscaling.

upvoted 3 times

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 App Service plan named APSPPlan1 set to the Basic B1 pricing tier. APSPPlan1 contains an App Service web app named WebApp1.

You plan to enable schedule-based autoscaling for APSPPlan1.

You need to minimize the cost of running WebApp1.

Solution: Scale out APSPPlan1.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Community vote distribution

B (100%)

 **Mattt** Highly Voted 7 months, 3 weeks ago

**Selected Answer: B**

No

Scaling out (increasing the instance count) does not minimize costs; it would actually increase them because more instances of the App Service Plan would be running. To minimize costs, you could scale down to a lower tier that still supports autoscaling (such as the Standard tier if autoscaling is required) rather than scaling out.

upvoted 8 times

 **fk1234** Most Recent 5 months, 3 weeks ago

**Selected Answer: B**

Scale out only increases number of VMs, doesn't solve auto scaling issue. Answer is NO

upvoted 1 times

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 App Service plan named APSPPlan1 set to the Basic B1 pricing tier. APSPPlan1 contains an App Service web app named WebApp1.

You plan to enable schedule-based autoscaling for APSPPlan1.

You need to minimize the cost of running WebApp1.

Solution: Scale up APSPPlan1 to the Premium V2 pricing tier.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: A**

 **Mattt** Highly Voted 7 months, 3 weeks ago

**Selected Answer: B**

No


Scaling up to the Premium V2 pricing tier would increase the cost, not minimize it. The Premium V2 tier is more expensive than the Basic and Standard tiers. Since the goal is to enable schedule-based autoscaling while minimizing costs, the Standard tier would be more cost-effective, as it supports autoscaling at a lower price point than the Premium tiers.

upvoted 7 times

 **fiercefreedom** Most Recent 8 months, 1 week ago


Switching to Standard would be a better solution in order to minimize costs.

upvoted 3 times

 **fc61c73** 8 months, 1 week ago


No, the solution does not meet the goal. Scaling up to the Premium V2 pricing tier will provide more features and better performance, but it will not minimize the cost. In fact, the Premium V2 tier is more expensive than the Basic tier. To enable schedule-based autoscaling and minimize costs, you should consider scaling up to the Standard tier, which supports autoscaling features and is more cost-effective than the Premium V2 tier.

upvoted 2 times

 **\_tharindus** 8 months, 2 weeks ago

Given answer is correct. Auto-scaling is only supported from Standard App Service Plans and onwards.

upvoted 2 times

 **lincio96** 8 months, 1 week ago

sure, but the goal is also to minimize the costs so as you said it starts from standard onwards. So there's a possibility to keep costs lower. On this basis I'd say NO

upvoted 2 times

## HOTSPOT

-

You have an Azure Functions app using the Consumption hosting plan for a company. The app contains the following functions:

Function Name	Trigger type
f1	HTTP
f2	Timer
f3	Azure Queues

You plan to enable dynamic concurrency on the app. The company requires that each function has its concurrency level managed separately.

You need to configure the app for dynamic concurrency.

Which file or function names should you use? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

## Requirement

File name

## File or Function Name

  
function.json  
host.json  
local.settings.json

Function name

  
f1  
f2  
f3

## Suggested Answer:

## Answer Area

## Requirement

File name

## File or Function Name

  
function.json  
**host.json**  
local.settings.json

Function name

  
f1  
f2  
**f3**

 **Mattt**  7 months, 3 weeks ago

host.json

f3

Azure Queue Trigger (f3):

This is the most logical choice for concurrency management. Queue-triggered functions handle messages from an Azure Queue and can process multiple messages concurrently, making it essential to manage concurrency to control throughput and avoid overloading the function app or downstream resources.



- HTTP Trigger (f1):

While concurrency can technically be managed for HTTP-triggered functions, it is less common because HTTP triggers usually handle requests one at a time. However, if you expect high traffic, you may still want to limit concurrency to prevent excessive parallel processing.

- Timer Trigger (f2):

Timer triggers are generally used for scheduled tasks and typically don't require concurrency management, as they trigger based on a specific schedule and are designed to run as single instances.

upvoted 5 times

  **examtopicsLogin123** 6 months, 3 weeks ago

Agree

"You can enable dynamic concurrency ... in your host.json file"

"Dynamic concurrency is currently only supported for the Azure Blob, Azure Queue, and Service Bus triggers"

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-concurrency#dynamic-concurrency>

upvoted 4 times

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 Container Apps app and disable ingress on the container app.

Users report that they are unable to access the container app. You investigate and observe that the app has scaled to 0 instances.

You need to resolve the issue with the container app.

Solution: Enable ingress, create an TCP scale rule, and apply the rule to the container app.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

*Community vote distribution*

B (100%)

 **Mattt**  7 months, 3 weeks ago

**Selected Answer: B**

No

Answer: Enable ingress and configure the minimum replicas to 1 for the container app  
upvoted 12 times



## DRAG DROP

-

You have two Azure Container Registry (ACR) instances: ACR01 and ACR02.

You plan to implement a containerized application named APP1 that will use a base image named BASE1. The image for APP1 will be stored in ACR01. The image BASE1 will be stored in ACR02.

You need to automate the planned implementation by using a sequence of five Azure command-line interface (Azure CLI) commands. Your solution must ensure that the APP1 image stored in ACR01 will be automatically updated when the BASE1 image is updated.

In which order should you perform the actions? To answer, move all container build automation options from the list of container build automations to the answer area and arrange them in the correct order.

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

## Container build automation

`az acr task credential add``az role assignment create``az acr build``az acr task run``az acr task create`

## Answer Area

1.

2.

3.


4.

5.

## Suggested Answer:

## Answer Area

1. `az acr build`2. `az acr task create`3. `az role assignment create`4. `az acr task credential add`5. `az acr task run`


 **NikolaMicic** Highly Voted 7 months, 2 weeks ago

The provided answer is correct:

`az acr build``az acr task create``az role assignment create``az acr task credential add``az acr task run`


<https://learn.microsoft.com/en-us/azure/container-registry/container-registry-tutorial-private-base-image-update>

upvoted 19 times

 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 2 times

  **abubabu** 7 months, 1 week ago

from gpt

az role assignment create:

Grant ACR01 access to ACR02 by creating a role assignment. This is necessary for ACR01 to pull the BASE1 image from ACR02 during the build process.

az acr task credential add:

Add credentials to ACR01 for authenticating with ACR02. This ensures the build task can access the BASE1 image in ACR02.

az acr task create:

Create a task in ACR01 that defines the build process for APP1, including the trigger for the BASE1 image in ACR02. This sets up the automation pipeline.

az acr task run:

Trigger the build task manually to ensure that the process is correctly configured and working as expected.

az acr build:

Use the build command to build the APP1 image in ACR01. This builds the APP1 image with the BASE1 image as the base.

upvoted 1 times

  **Dan696969** 6 months, 1 week ago

if you feed chatgpt the link from @NikolaMicic

chatgpt says this

az acr build



az acr task create

az role assignment create

az acr task credential add

az acr task run

upvoted 1 times

  **Jay456** 7 months, 1 week ago

What is automated about manually running these commands? :-o

upvoted 1 times

  **tndias** 7 months, 3 weeks ago

1- az acr task create

2- az role assignment create

3- az acr task credential add

4- az acr build

5- az acr task run

So, the correct order is to build the image first (az acr build) and then run the task (az acr task run). This ensures that the image is available in your Azure Container Registry before running any tasks that depend on

upvoted 3 times

  **Mattt** 7 months, 3 weeks ago

1- az acr task create

2- az role assignment create

3- az acr task credential add

4- az acr task run

5- az acr build

- Create ACR Task: This defines the build process for APP1, including the base image, build context, and output image.

- Assign Role: This ensures the task has the necessary permissions to access both registries.

- Add Credentials: This provides the task with authentication information to pull the BASE1 image and push the APP1 image.

- Run Task: This triggers the build process, which will automatically pull the latest BASE1 image and build APP1.
  - Manual Build: This is an alternative approach, but it requires manual intervention to trigger the build whenever BASE1 changes.
- upvoted 1 times

HOTSPOT

-

You plan to implement an Azure function named Function1 that will use the timer trigger.

You plan to use a TimeSpan value to set the schedule of function execution.

You need to select the hosting model and the schedule expression assignment method.

Which hosting model and schedule expression assignment method should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Function configuration

#### Function configuration option

#### Value

Hosting model

Consumption  
App Service plan  
Azure Arc-enabled Kubernetes

Schedule expression assignment

App settings  
host.json file  
function.json file

Suggested Answer:

#### Function configuration option

#### Value

Hosting model

Consumption  
App Service plan  
Azure Arc-enabled Kubernetes

Schedule expression assignment

App settings  
host.json file  
function.json file

 **21c5438** Highly Voted 5 months, 2 weeks ago

App Service plan

App settings

"A TimeSpan can be used only for a function app that runs on an App Service Plan. You can put the schedule expression in an app setting and set this property to the app setting name wrapped in % signs, as %ScheduleAppSetting%."

reference: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-timer?tabs=python-v2%2Cisolated-process%2Cnodejs-v4&pivots=programming-language-csharp>

upvoted 14 times

 **21c5438** 5 months, 2 weeks ago

Reference correction:

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-timer?tabs=python-v2%2Cisolated-process%2Cnodejs-v4&pivots=programming-language-csharp#attributes>



upvoted 4 times

  **vizay** Most Recent 2 months ago

Consumption

Function.json



upvoted 1 times

  **BenjaminBao** 3 months, 1 week ago

Consumption

Function.json

upvoted 1 times

  **e66b406** 5 months, 2 weeks ago

A CRON expression or a TimeSpan value. A TimeSpan can be used only for a function app that runs on an App Service Plan. You can put the schedule expression in an app setting and set this property to the app setting name wrapped in % signs, as %ScheduleAppSetting%.

upvoted 3 times

## HOTSPOT

-

You are developing a microservices-based application that uses Azure Container Apps. The application consists of several containerized services that handle tasks, such as processing orders, managing inventory, and generating reports.

You must secure the container apps. All apps must reside in the same virtual network, share the same Dapr configuration, and share the same logging location.

Apps must support the configuration of the amount of memory and compute resources available to containers.

You need to configure the Azure Container App.

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

NOTE: Each correct selection is worth one point.

**Azure Container App CLI command**

`az containerapp`

env	▼
create	
add-on	

env	▼
create	
add-on	

`-n MyContainerApp -g MyResourceGroup --location eastus2`

▼
--enable-mtls
--internal-only
--enable-workload-profiles

**Suggested Answer:**

`az containerapp`

env	▼
create	
add-on	

env	▼
create	
add-on	

`-n MyContainerApp -g MyResourceGroup --location eastus2`

▼
--enable-mtls
--internal-only
--enable-workload-profiles

21c5438 Highly Voted 5 months, 2 weeks ago

Provided answer is correct.

References:

<https://learn.microsoft.com/en-us/cli/azure/containerapp/env?view=azure-cli-latest#az-containerapp-env-create>

<https://learn.microsoft.com/en-us/azure/container-apps/workload-profiles-overview>

upvoted 5 times

DamuKeesh Most Recent 5 months ago

--internal-only: Ensures all apps reside in the same virtual network (VNet), meeting the networking requirement.

--enable-workload-profiles: Allows configuration of memory and compute resources for containers, addressing the resource management

requirement.



So these 2 option looks good. Idk which one to choose ?

upvoted 1 times

  **DamuKeesh** 5 months ago

It seems wrong answer

upvoted 1 times

  **wafa\_chaari** 5 months, 1 week ago

correct-

az containerapp env create : to Create a Container Apps environment.

- workload profile determines the amount of compute and memory resources available

upvoted 1 times

## HOTSPOT -

You are developing a solution that uses the Azure Storage Client library for .NET. You have the following code: (Line numbers are included for reference only.)

```

01 CloudBlockBlob src = null;
02 try
03 {
04     src = container.ListBlobs().OfType<CloudBlockBlob>().FirstOrDefault();
05     var id = await src.AcquireLeaseAsync(null);
06     var dst = container.GetBlockBlobReference(src.Name);
07     string cpid = await dst.StartCopyAsync(src);
08     await dst.FetchAttributeAsync();
09     return id;
10 }
11 catch (Exception e)
12 {
13     throw;
14 }
15 finally
16 {
17     if (src != null)
18         await src.FetchAttributesAsync();
19     if (src.Properties.LeaseState != LeaseState.Available)
20         await src.BreakLeaseAsync(new TimeSpan(0));
21 }

```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statement	Yes	No
The code creates an infinite lease	<input type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input type="radio"/>
The finally block releases the lease	<input type="radio"/>	<input type="radio"/>

### Answer Area

Suggested Answer:

Statement	Yes	No
The code creates an infinite lease	<input checked="" type="radio"/>	<input type="radio"/>
The code at line 06 always creates a new blob	<input type="radio"/>	<input checked="" type="radio"/>
The finally block releases the lease	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes -

AcquireLeaseAsync does not specify leaseTime.

leaseTime is a TimeSpan representing the span of time for which to acquire the lease, which will be rounded down to seconds. If null, an infinite lease will be acquired. If not null, this must be 15 to 60 seconds.

Box 2: No -

The GetBlockBlobReference method just gets a reference to a block blob in this container.

Box 3: Yes -

The BreakLeaseAsync method initiates an asynchronous operation that breaks the current lease on this container.

Reference:



<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync>  
<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.getblockblobreference>  
<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync>

🗨️ 👤 **Tealon** Highly Voted 4 years, 6 months ago

I think the answer is correct:

Optional. Version 2012-02-12 and newer. For a break operation, this is the proposed duration of seconds that the lease should continue before it is broken, between 0 and 60 seconds. This break period is only used if it is shorter than the time remaining on the lease. If longer, the time remaining on the lease is used. A new lease will not be available before the break period has expired, but the lease may be held for longer than the break period. If this header does not appear with a break operation, a fixed-duration lease breaks after the remaining lease period elapses, and an infinite lease breaks immediately.

From: <https://docs.microsoft.com/en-us/rest/api/storageservices/lease-blob>  
upvoted 47 times

🗨️ 👤 **AzureFailure** 4 years ago

It talks about the header values only ReleaseLeaseAsync release lease  
upvoted 3 times

🗨️ 👤 **AzureAz204Fan** 4 years, 1 month ago

Last few words "infinite lease breaks immediately" is the key for this context.  
upvoted 4 times

🗨️ 👤 **Figa** Highly Voted 4 years, 7 months ago

I think Box 3 should be no instead of yes. BreakRelease doesn't release the lease directly.

You use ReleaseLease to do this:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.releaselease?view=azure-dotnet-legacy>  
upvoted 24 times

🗨️ 👤 **Figa** 4 years, 7 months ago

Ah I think model is right:  
BreakLeaseAsync (TimeSpan? breakPeriod)

breakPeriod

Nullable<TimeSpan>

A TimeSpan representing the amount of time to allow the lease to remain, which will be rounded down to seconds. If null, the break period is the remainder of the current lease, or zero for infinite leases.

upvoted 4 times

🗨️ 👤 **Chang401** 2 years, 7 months ago

ya!!! an infinite lease breaks immediately  
upvoted 1 times

🗨️ 👤 **wafa\_chaari** Most Recent 5 months, 1 week ago

correct . yes no yes - acquire leaseasync(null)--> infinite , On line 6 it gets a reference to a block blob in this container.

A lease is a lock on a blob or file for write and delete operations<sup>12</sup>. The lock duration can be 15 to 60 seconds, or can be infinite  
upvoted 1 times

🗨️ 👤 **4bd3116** 9 months, 2 weeks ago

The code creates an infinite lease:

The AcquireLeaseAsync method is called with null as the parameter, which does not create an infinite lease. To create an infinite lease, it should be called with -1.

Answer: No

The code at line 06 always creates a new blob:

Line 06 (var dst = container.GetBlockBlobReference(src.Name);) gets a reference to an existing blob with the same name as src. It does not create a new blob but rather references an existing one.

Answer: No

The finally block releases the lease:

The finally block contains await src.BreakLeaseAsync(new TimeSpan(0)), which releases the lease on the blob.

Answer: Yes

Thanks for pointing that out! Is there anything else you'd like to discuss or clarify?

upvoted 4 times

  **4bd3116** 10 months ago

Lease Duration:

No, the code does not create an infinite lease. The lease acquired at line 5 (`await src.AcquireLeaseAsync(null)`) does not specify a duration (the null argument). By default, this results in a finite lease.

If you want to create an infinite lease, you can provide a non-null `TimeSpan` argument to `AcquireLeaseAsync`, specifying the desired lease duration.

Creating a New Blob:

No, the code at line 6 (`var dst = container.GetBlockBlobReference(src.Name);`) does not always create a new blob. Instead, it creates a new `CloudBlockBlob` reference named `dst` that points to the same blob as `src`.

Essentially, `dst` refers to the blob with the same name as `src`.

Releasing the Lease:

Yes, the final code ensures that the lease is released.

In the finally block (lines 15-25), it checks if `src` is not null:

Fetches attributes of `src`.


If the lease state of `src` is not available (meaning it's still leased), it breaks the lease using `src.BreakLeaseAsync(new TimeSpan(0))`.

upvoted 1 times

  **CafeLago** 10 months, 3 weeks ago

Finally block doesn't release the lease. Is nobody else able to see the if above the `BreakLeaseAsync`?

upvoted 1 times

  **CafeLago** 10 months, 3 weeks ago

Sorry, take it back. `LeaseState.Available` means not `LeaseState.Lease`d, so not `Available` means `Lease`d, and so the lease can be broken.

upvoted 2 times

  **jobolesonihal** 1 year, 3 months ago

Answer is Correct.

`BlobLease.BreakLeaseAsync`:

If the break period is not specified, the default value is 0 (immediate break).

upvoted 1 times

  **bgbgvfvf** 1 year, 6 months ago



I think the answer is correct

upvoted 2 times

  **dom271219** 1 year, 9 months ago



So, when you provide a `TimeSpan` value of 0, it essentially means "break the lease right now without any delay." This behavior aligns with the statement you provided: "If null, the break period is the remainder of the current lease, or zero for infinite leases." When you explicitly set it to 0, it ensures an immediate lease termination. It is a matter of knowing the English Language.

upvoted 2 times

  **kabbas** 1 year, 11 months ago

to break the lease it will have to enter in the if block ..and since the lease `state` will be always be available since it is infinite ...it will never enter that block...so the answer is no

upvoted 1 times

  **Esward** 2 years, 5 months ago

Yes, No, Yes

upvoted 1 times

  **Esward** 2 years, 5 months ago

`AcquireLease(null)`-->Creates infinite lease

upvoted 2 times

  **Esward** 2 years, 5 months ago

`AcquireLease(-1)` will also create infinite lease

upvoted 1 times

  **coffecold** 2 years, 8 months ago

version 11 is legacy.

Similar code for version 12 see

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-copy?tabs=dotnet>

upvoted 1 times

🗨️ 👤 **coffecold** 2 years, 8 months ago

This is Microsoft.Azure.Storage.Blob v11.1.0 code, it is a legacy library. Did anyone got this on the exam recently?

upvoted 2 times

🗨️ 👤 **strongmmc** 2 years, 9 months ago

i'm looking for an explanation on what a "lease" is, or what is intended by it, and cannot find any explanatory resource. Does anyone have a reference to it ?

upvoted 3 times

🗨️ 👤 **strongmmc** 2 years, 9 months ago

<https://www.c-sharpcorner.com/article/implementing-blob-leasing-understanding-blob-storage-part-9/>

The Lease Blob operation establishes and manages a lock on a blob for write and delete operations. The lock duration can be 15 to 60 seconds or it can be infinite. Once you have the lease, you can update the Blob or delete the Blob without worrying about another process changing it underneath you. When a Blob is leased, other processes can still read it but any attempt to update it will fail. You can also update Blobs without taking a lease first, but you do run the chance of another process also attempting to modify it at the same time.

upvoted 4 times

🗨️ 👤 **winston\_45** 2 years, 10 months ago

Correct answer:

YES

NO

YES

upvoted 3 times

🗨️ 👤 **gds03** 2 years, 10 months ago

If the finally code releases the lease, then 1st is wrong. The code doesn't create an infinite lease. Answer is NO.

upvoted 4 times

🗨️ 👤 **Emil\_Topics** 1 year, 7 months ago

I agree. It's a bit silly to create a lease and then removing it, but that is in my opinion what is happening here. Why do people agree that the first one should be YES?

upvoted 1 times

🗨️ 👤 **sasha\_gud** 3 years, 3 months ago

Tried to execute the code sample:

- 1) If acquires lease on the blob with null parameter, means the infinite lease is acquired [1]
- 2) On line 6 it gets reference to existing source blob, so no new blob is created there
- 3) It fails on start copy operation, because it tries to copy onto itself and the lease is already acquired
- 4) Goes to catch and then finally block
- 5) In finally breaks the lease with (Zero) parameter that means the immediate lease break [2]

[1] [https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync?view=azure-dotnet-legacy#microsoft-azure-storage-blob-cloudblobcontainer-acquireleaseasync\(system-nullable\(\(system-timespan\)\)-system-string\)](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.acquireleaseasync?view=azure-dotnet-legacy#microsoft-azure-storage-blob-cloudblobcontainer-acquireleaseasync(system-nullable((system-timespan))-system-string))

[2] [https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync?view=azure-dotnet-legacy#microsoft-azure-storage-blob-cloudblobcontainer-breakleaseasync\(system-nullable\(\(system-timespan\)\)\)](https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.blob.cloudblobcontainer.breakleaseasync?view=azure-dotnet-legacy#microsoft-azure-storage-blob-cloudblobcontainer-breakleaseasync(system-nullable((system-timespan))))

upvoted 16 times

You are building a website that uses Azure Blob storage for data storage. You configure Azure Blob storage lifecycle to move all blobs to the archive tier after 30 days.

Customers have requested a service-level agreement (SLA) for viewing data older than 30 days.

You need to document the minimum SLA for data recovery.

Which SLA should you use?

- A. at least two days
- B. between one and 15 hours
- C. at least one day
- D. between zero and 60 minutes


**Suggested Answer: B**

The archive access tier has the lowest storage cost. But it has higher data retrieval costs compared to the hot and cool tiers. Data in the archive tier can take several hours to retrieve depending on the priority of the rehydration. For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal>

Community vote distribution



🗳️ **homimi6115** Highly Voted 4 years, 7 months ago

Answer is correct

upvoted 45 times

🗳️ **surprise0011** 2 years, 2 months ago

received 2023-04-17 went given answer, score 926

upvoted 5 times

🗳️ **130nk3r5** 1 year, 6 months ago

Got this today.

Went with answer here.

Score 927

upvoted 1 times

🗳️ **mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: B

- Standard priority: The rehydration request will be processed in the order it was received and may take up to 15 hours.

- High priority: The rehydration request will be prioritized over Standard requests and may finish in under 1 hour for objects under ten GB in size.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers?tabs=azure-portal#archive-access-tier>

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-rehydration?tabs=azure-portal>

upvoted 27 times

🗳️ **wafa\_chaari** Most Recent 5 months, 1 week ago

Selected Answer: B

Data in the archive tier can take up to 15 hours to rehydrate

upvoted 2 times

🗳️ **Vichu\_1607** 7 months, 3 weeks ago

Selected Answer: B

B. between one and 15 hours

Explanation:

Archive Tier Retrieval Time: The retrieval time for data stored in the archive tier typically ranges from one hour to 15 hours. This is the time it takes to rehydrate the data from the archive tier to an online tier (hot or cool) where it can be accessed.

Service-Level Agreement (SLA): When documenting the SLA for data recovery, you should account for the maximum retrieval time. Therefore, the minimum SLA for data recovery from the archive tier should be between one and 15 hours.

upvoted 2 times

🗳️ 👤 **IntegrationTeam** 10 months, 1 week ago

**Selected Answer: A**

Retrieving from the Archive tier takes 12 to 48 hours

upvoted 1 times

🗳️ 👤 **manopeydakon** 1 year, 5 months ago

The SLA for data recovery from the archive tier in Azure Blob storage is indeed between zero and 60 minutes. Option B ("between one and 15 hours") is not accurate in the context of Azure Blob storage archive tier recovery. The correct answer should be:

D. between zero and 60 minutes

upvoted 1 times

🗳️ 👤 **JH81** 2 years ago

**Selected Answer: B**

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 1 times

🗳️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

Easy one.

My case:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 1 times

🗳️ 👤 **methkill** 1 year, 7 months ago

These are some random bot answers.

upvoted 2 times

🗳️ 👤 **Saluk\_DE** 2 years, 3 months ago

Question was in Exam 2023-03-30

upvoted 1 times

🗳️ 👤 **sarmaria** 2 years, 3 months ago

Got this question in the exam on 16/03/2023. Went with proposed solution. Make sure to prepare for case studies. I got city and lights case study.

upvoted 2 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

B. between one and 15 hour

upvoted 1 times

🗳️ 👤 **elequiel** 2 years, 7 months ago

**Selected Answer: B**

Correct Answer: B,

upvoted 1 times

🗳️ 👤 **winston\_45** 2 years, 10 months ago

This question is sooo open for interpretation: "For small objects, a high priority rehydrate may retrieve the object from archive in under 1 hour."

Where is this 15 hours coming from?

upvoted 4 times

🗳️ 👤 **AZAdam22** 2 years, 11 months ago

**Selected Answer: B**

B - Because it takes between one and 15 hours to recover data from the archive tier.

upvoted 2 times

🗨️ 👤 **Eltooth** 3 years ago

**Selected Answer: B**

B is correct answer.

Gen 1 storage used to be up to 24 hours, now gen 2 is up to 15 hours.

<https://docs.microsoft.com/en-us/azure/storage/blobs/archive-rehydrate-overview#rehydration-priority>

upvoted 2 times

🗨️ 👤 **meoukg** 3 years, 3 months ago

Got it on 03/2022, I chose B. between one and 15 hours

upvoted 1 times

🗨️ 👤 **KiranAtShinde** 3 years, 5 months ago

**Selected Answer: B**

Correct Answer is between 1-15 hours

upvoted 1 times

## HOTSPOT -

You are developing a ticket reservation system for an airline.

The storage solution for the application must meet the following requirements:

- ⇒ Ensure at least 99.99% availability and provide low latency.
- ⇒ Accept reservations even when localized network outages or other unforeseen failures occur.
- ⇒ Process reservations in the exact sequence as reservations are submitted to minimize overbooking or selling the same seat to multiple travelers.
- ⇒ Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window.

You provision a resource group named `airlineResourceGroup` in the Azure South-Central US region.

You need to provision a SQL API Cosmos DB account to support the app.

How should you complete the Azure CLI commands? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
resourceGroupName= 'airlineResourceGroup'
name= 'docdb-airline-reservations'
databaseName= 'docdb-tickets-database'
collectionName= 'docdb-tickets-collection'
consistencyLevel=
```

▼
Strong
Eventual
ConsistentPrefix
BoundedStaleness

```
az cosmosdb create \
--name $name \
```

▼
--enable-virtual-network true \
--enable-automatic-failover true \
--kind 'GlobalDocumentDB' \
--kind 'MongoDB' \

```
--resource-group $resourceGroupName \
--max-interval 5 \
```

▼
--locations 'southcentralus'
--locations 'eastus'
--locations 'southcentralus=0 eastus=1 westus=2'
--locations 'southcentralus=0'

```
--default-consistency-level = $consistencylevel
```

## Answer Area

```
resourceGroupName= 'airlineResourceGroup'
name= 'docdb-airline-reservations'
databaseName= 'docdb-tickets-database'
collectionName= 'docdb-tickets-collection'
consistencyLevel=
```

Strong
Eventual
ConsistentPrefix
BoundedStaleness

```
az cosmosdb create \
--name $name \
```

Suggested Answer:

```
--enable-virtual-network true \
--enable-automatic-failover true \
--kind 'GlobalDocumentDB' \
--kind 'MongoDB' \
--resource-group $resourceGroupName \
--max-interval 5 \
--locations 'southcentralus'
--locations 'eastus'
--locations 'southcentralus=0 eastus=1 westus=2'
--locations 'southcentralus=0'
--default-consistency-level = $consistencylevel
```

Box 1: BoundedStaleness -

Bounded staleness: The reads are guaranteed to honor the consistent-prefix guarantee. The reads might lag behind writes by at most "K" versions (that is, "updates") of an item or by "T" time interval. In other words, when you choose bounded staleness, the "staleness" can be configured in two ways:

The number of versions (K) of the item

The time interval (T) by which the reads might lag behind the writes

Incorrect Answers:

Strong -

Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Box 2: --enable-automatic-failover true\

For multi-region Cosmos accounts that are configured with a single-write region, enable automatic-failover by using Azure CLI or Azure portal. After you enable automatic failover, whenever there is a regional disaster, Cosmos DB will automatically failover your account.

Question: Accept reservations event when localized network outages or other unforeseen failures occur.

Box 3: --locations'southcentralus=0 eastus=1 westus=2

Need multi-region.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/manage-with-cli.md>

 **sasisang**  4 years, 7 months ago

Answer is correct.

upvoted 73 times

 **cbn** 4 years, 4 months ago

Last option is correct as well

<https://stackoverflow.com/questions/51197375/creating-cosmosdb-with-azure-cli-in-powershell>

upvoted 7 times

 **cbn** 4 years, 4 months ago



I guess there is a typo on quotes though.

upvoted 5 times

  **ning**  3 years, 11 months ago

No need to overthink here, max-interval, indicates this must be bounded-staleness, enable-automatic-failover, indicated this must be multi-region  
upvoted 65 times

  **OPT\_001122** 2 years, 6 months ago


this is best explanation  
upvoted 2 times

  **laminall**  2 months ago



```
az cosmosdb create \  
--name airlineReservationDB \  
--resource-group airlineResourceGroup \  
--locations regionName=southcentralus failoverPriority=0 isZoneRedundant=true \  
--default-consistency-level BoundedStaleness \  
--enable-automatic-failover true \  
--enable-multiple-write-locations true \  
--capabilities EnableServerless \  
--max-staleness-prefix 100 \  
--max-interval 5  
upvoted 1 times
```

  **examtopics44556677** 1 year, 3 months ago



Confused how most of the people here scored same score - 927.  
I see 927 everywhere.  
upvoted 4 times

  **Occ50bf** 10 months, 3 weeks ago

Bots, buddy. They're bots.  
upvoted 2 times

  **jakobste** 1 year, 10 months ago



It's important to carefully read the entire question.  
upvoted 3 times

  **Dianahu** 1 year, 11 months ago



I think that enable-automatic-failover = true option does not make any sense without other regions specified but given syntax for regions and priorities is wrong and does not work  
upvoted 1 times

  **Dats1987** 2 years, 2 months ago



I think ans should be using chat gpt:  
1. Eventual : maintain high availability and low latency  
2. enable-automatic-failover = true  
3. locations "South Central US"=0 : This is primary region  
upvoted 2 times

  **Ivanlco** 1 year, 12 months ago

Eventual is definitely not correct  
upvoted 10 times

  **AzSimuz** 2 years, 3 months ago

BoundedStaleness doesn't guarantee the order of the order, that's session consistency. Where am I wrong?  
upvoted 2 times

  **Esward** 2 years, 4 months ago

It was there in 13 Feb 2023 exam  
upvoted 5 times

  **mdg3501** 2 years, 6 months ago

got this on 2022-12-7  
upvoted 2 times

  **leandrocantiero** 3 years, 3 months ago

esse aqui ta tudo errado

upvoted 4 times

🗨️ 👤 **glam** 4 years, 1 month ago

correct

upvoted 2 times

🗨️ 👤 **crepatata** 4 years, 2 months ago

The strong consistency doesn't work here:

1) there is a `--max-interval` property which is being used with bounded staleness only

2) there is a requirement "Accept reservations event when localized network outages or other unforeseen failures occur." which points us to multiple writes for multiple regions which is not being supported by the Strong consistency.

`--locations` syntax looks like obsolete as for late march 2021

upvoted 11 times

🗨️ 👤 **Spooky7** 4 years, 1 month ago

Ad. 2 - it is not true that it implies writes to multiple regions. In this case there is single write region with automatic failover configured (which provides write availability). Besides with write to multiple regions BoundedStaleness minimum time is 300 seconds, while requirement is 5 seconds (which is actually a minimum for BoundedStaleness with single write region)

upvoted 3 times

🗨️ 👤 **Ralph** 4 years, 2 months ago

Answer is wrong .. "strong" instead of "boundedStaleness"

upvoted 4 times

🗨️ 👤 **TakumaK** 4 years, 1 month ago

Are you sure? If then, elaborate why. Does this option `--max-interval` work for strong?

upvoted 8 times

🗨️ 👤 **[Removed]** 4 years, 5 months ago

I think this should be "locations southcentralus". The requirement for 99.99% availability is covered by a single region and the syntax for providing multiple regions is wrong.

<https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability>

[https://docs.microsoft.com/en-us/cli/azure/cosmosdb?view=azure-cli-latest#az\\_cosmosdb\\_check\\_name\\_exists-optional-parameters](https://docs.microsoft.com/en-us/cli/azure/cosmosdb?view=azure-cli-latest#az_cosmosdb_check_name_exists-optional-parameters)

upvoted 3 times

🗨️ 👤 **agcertif** 4 years, 5 months ago

I agree with you.... 99,99 can be reached with only one region.

Why it's not the right solution (someone can helps please) ?

upvoted 3 times

🗨️ 👤 **Brak** 4 years, 5 months ago

Because of the need for high availability during local network outages. Hence the auto failover selection.

upvoted 5 times

🗨️ 👤 **NikB** 4 years, 4 months ago

Agreed, also according to the documentation, multiple regions have a min T of 300s:

"For a single region account, the minimum value of K and T is 10 write operations or 5 seconds. For multi-region accounts the minimum value of K and T is 100,000 write operations or 300 seconds."

upvoted 2 times

🗨️ 👤 **RaviKS** 4 years, 5 months ago

Given Answer is Correct

upvoted 4 times

🗨️ 👤 **tomis** 4 years, 5 months ago

`--locations` : Add a location to the Cosmos DB database account.

Usage: `--locations KEY=VALUE [KEY=VALUE ...]`

Required Keys: `regionName`, `failoverPriority`

Optional Key: `isZoneRedundant`

Default: single region account in the location of the specified resource group.

Failover priority values are 0 for write regions and greater than 0 for read regions. A

failover priority value must be unique and less than the total number of regions.

Multiple locations can be specified by using more than one `--locations`` argument.

upvoted 8 times

  **diligent176** 4 years, 5 months ago

You are correct about this - none of the given options will work (all wrong syntax).

But I think they "meant" to use the 3 regions answer.

upvoted 4 times

## HOTSPOT -

You are preparing to deploy a Python website to an Azure Web App using a container. The solution will use multiple containers in the same container group. The

Dockerfile that builds the container is as follows:

```
FROM python:3
ADD website.py
CMD [ "python", "./website.py"]
```

You build a container by using the following command. The Azure Container Registry instance named images is a private registry.

```
docker build -t images.azurecr.io/website:v1.0.0
```

The user name and password for the registry is admin.

The Web App must always run the same version of the website regardless of future builds.

You need to create an Azure Web App to run the website.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

```
az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan
```

	▼
--sku SHARED	
--tags container	
--sku B1 --hyper-v	
--sku B1 --is-linux	

```
az webapp create --plan websitePlan
```

	▼
--deployment-source-url images.azurecr.io/website:v1.0.0	
--deployment-source-url images.azurecr.io/website:latest	
--deployment-container-image-name images.azurecr.io/website:v1.0.0	
--deployment-container-image-name images.azurecr.io/website:latest	

```
az webapp config
```

	▼
set --python-version 2.7 --generic-configurations user=admin password=admin	
set --python-version 3.6 --generic-configurations user=admin password=admin	
container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin	
container set --docker-registry-server-url https://images.azurecr.io/website -u admin -p admin	

Suggested Answer:

## Answer Area

```
az configure --defaults web=website
az configure --defaults group=website
az appservice plan create --name websitePlan
```

--sku SHARED
--tags container
--sku B1 --hyper-v
--sku B1 --is-linux

```
az webapp create --plan websitePlan
```

--deployment-source-url images.azurecr.io/website:v1.0.0
--deployment-source-url images.azurecr.io/website:latest
--deployment-container-image-name images.azurecr.io/website:v1.0.0
--deployment-container-image-name images.azurecr.io/website:latest

```
az webapp config
```

set --python-version 2.7 --generic-configurations user=admin password=admin
set --python-version 3.6 --generic-configurations user=admin password=admin
container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin
container set --docker-registry-server-url https://images.azurecr.io/website -u admin -p admin

Box 1: --SKU B1 --hyper-v -

--hyper-v

Host web app on Windows container.

Box 2: --deployment-source-url images.azurecr.io/website:v1.0.0

--deployment-source-url -u

Git repository URL to link with manual integration.

The Web App must always run the same version of the website regardless of future builds.

Incorrect:

--deployment-container-image-name -i

Linux only. Container image name from Docker Hub, e.g. publisher/image-name:tag.

Box 3: az webapp config container set -url https://images.azurecr.io -u admin -p admin az webapp config container set

Set a web app container's settings.

Parameter: --docker-registry-server-url -r

The container registry server url.

The Azure Container Registry instance named images is a private registry.

Example:

az webapp config container set --docker-registry-server-url https://{azure-container-registry-name}.azurecr.io

Reference:

<https://docs.microsoft.com/en-us/cli/azure/appservice/plan>

 **glam**  3 years, 7 months ago

--sku B1 --is-linux

--deployment-container-image-name images.azurecr.io/website:v1.0.0

-- container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin

upvoted 160 times

 **Esward** 1 year, 11 months ago

[https://www.linkedin.com/pulse/how-quickly-create-micro-service-azure-webapp-fastapi-bonnet?trk=pulse-article\\_more-articles\\_related-content-card](https://www.linkedin.com/pulse/how-quickly-create-micro-service-azure-webapp-fastapi-bonnet?trk=pulse-article_more-articles_related-content-card)

upvoted 2 times

 **applepie** 1 year, 5 months ago

got this question today, go with this answer - 7/30/2023, score 895/1000



upvoted 4 times

 **aperez1979**  3 years, 8 months ago

"use multiple containers in the same container group" this not is possible in windows.

Solution is:

--is-linux  
--deployment-container-image-name  
upvoted 46 times

  **Tom87** 3 years, 8 months ago

That's true.

"Multi-container groups are currently restricted to Linux containers."

<https://docs.microsoft.com/en-us/azure/container-instances/container-instances-multi-container-group>

upvoted 11 times

  **azurelearner666** 3 years, 6 months ago

True! also last response is correct with

-- container set --docker-registry-server-url <https://images.azurecr.io> -u admin -p admin

upvoted 5 times

  **azureq9** Most Recent 7 months, 2 weeks ago

Box 1 -- --sku B1 --is-linux. You need an application service plan that supports Linux because your Docker container is based on Linux images (python:3). SKU B1 is the basic level and can run on Linux, which should suit your requirements.

Box 2 -- --deployment-container-image-name images.azurecr.io/website:v1.0.0. You need to specify a specific container image name to ensure that the Web App always runs the same version of the website. Using the v1.0.0 tag ensures that this version will not change when new builds are released.

Box 3 -- container set --docker-registry-server-url <https://images.azurecr.io> -u admin -p admin. You need to configure Docker registry credentials so that Web App can pull images from the Azure Container Registry. This selection provides the correct URL of the registry and the admin login information you provided.



upvoted 8 times

  **neelkanths** 8 months ago

Got it on 20 April 2024...Marks > 900...answer is correct....all questions from examtopics 400 question bank



Answer is correct

upvoted 1 times

  **1989jay** 11 months, 3 weeks ago

This question was on exam. 09/01/2024. Passed 872.

upvoted 2 times

  **leviatas** 1 year, 3 months ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 1 times

  **NightshadeRC** 1 year, 5 months ago

Had this question today: 2023-07-26

upvoted 1 times

  **surprise0011** 1 year, 8 months ago



microsoft promote good practises for -u and -p (◡‿◡)

upvoted 9 times

  **[Removed]** 1 year, 8 months ago



Got this in exam today (5 April 2023)

upvoted 5 times

  **uffuchsi** 1 year, 10 months ago



Received this in my exam today (22/02/2023). Selected linux, --deployment-container...v1.0.0, and container set ... in --u admin --p admin. Score 927.

upvoted 5 times

  **Esward** 1 year, 10 months ago



It was there in 13 Feb 2023 exam

upvoted 4 times

  **Tito28** 1 year, 12 months ago


this answer are not wrong, if you must know, these question and answer are giving by microsoft , so how can it be wrong if microsoft support those answer?

upvoted 1 times

  **Net\_IT** 1 year, 10 months ago



The answers are not provided by Microsoft but by the community/Examtopics, sometimes given answers are wrong. But the discussions are great to help you with that! In this case I believe the highest voted answers are correct (from aperez1979).

upvoted 5 times

  **Dani\_ac7** 2 years, 5 months ago



Linux for multiple containers

upvoted 3 times

  **Pize** 2 years, 6 months ago

Another wrong answer - this site riddled with wrong answers

upvoted 7 times

  **Eltooth** 2 years, 6 months ago

--sku B1 --is-linux

--deployment-container-image-name images.azurecr.io/website:v1.0.0

-- container set --docker-registry-server-url https://images.azurecr.io -u admin -p admin

upvoted 4 times

  **yevgen91** 3 years ago

The suggested answer is for Windows, but Python is deprecated for Windows, as well as multiple containers are not possible in Windows.

So it must be Linux, and based on this article <https://docs.microsoft.com/en-us/azure/app-service/tutorial-custom-container?pivots=container-linux>

the answer must be:

1. --sku B1 --is-linux

2. --deployment-container-image-name ... :v1.0.0

3. container set ... https://images.azurecr.io ...

upvoted 7 times

  **Molte** 3 years ago

for answer 3: Can somebody tell me why its images.azurecr.io and not images.azurecr.io/website?

upvoted 4 times

  **ScubaDiver123456** 2 years, 11 months ago

The command requires the URL of the registry, not the image name. That is appended later.

"--docker-registry-server-url

The container registry server url."

upvoted 9 times

## HOTSPOT -

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

A rule already exists to scale up the App Service when the average queue length of unprocessed and valid queue messages is greater than 1000.

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

How should you configure the Scale rule? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**Scale rule ×

## Metric source

Storage queue
Service Bus queue
Current resource
Storage queue (classic)

## Resource type

Service Bus Namespaces
------------------------

## Resource

MessageQueue1103
------------------

## \* Queues

itemqueue
-----------

## Criteria

## \* Metric name

Message Count
Active Message Count

1 minute time grain

## \* Time grain statistic

Total
Maximum
Average
Count

## \* Operator

--



## Answer Area

### Scale rule

#### Metric source

Storage queue  
Service Bus queue  
Current resource  
Storage queue (classic)

#### Resource type

Service Bus Namespaces

#### Resource

MessageQueue1103

#### \* Queues

itemqueue

#### Criteria

##### \* Metric name

Message Count  
Active Message Count

1 minute time grain

##### \* Time grain statistic

Total  
Maximum  
Average  
Count

##### \* Operator

Suggested Answer:

Box 1: Service bus queue -

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

Box 2: ActiveMessage Count -

ActiveMessageCount: Messages in the queue or subscription that are in the active state and ready for delivery.

Box 3: Count -

Box 4: Less than or equal to -

You need to add a new rule that will continuously scale down the App Service as long as the scale up condition is not met.

Box 5: Decrease count by

 **VijiVivek**  3 years, 8 months ago

The correct answers are

- 1) Service bus queue
- 2) Active message count
- 3) Average
- 4) Less than or equal to
- 5) Decrease count by

upvoted 129 times

 **micro9000** 2 years, 1 month ago

I think the #4 is should be "Less than" only, because you don't want to trigger the Scale up and Scale down at the same time.

upvoted 7 times

 **JustHereToLearn** 1 year, 11 months ago

I won't trigger at same time as question says "average queue length of unprocessed and valid queue messages is greater than 1000"

upvoted 4 times

 **micro9000** 2 years, 1 month ago

Here is a reference why I think it should be Less than,

<https://learn.microsoft.com/en-us/training/modules/scale-apps-app-service/5-autoscale-best-practices?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wvl.create-azure-app-service-web-apps>

upvoted 1 times

  **azurelearner666** 3 years, 6 months ago

this is correct (the "official" solution is wrong)

Also check the proper image [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_062\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_062_Image_0001.jpg)

upvoted 42 times

  **Evo\_Morales** 2 years, 8 months ago




How often are the responses scrubbed for updates/corrections?

upvoted 2 times

  **applepie** 1 year, 5 months ago

got this question today, go with this answer - 7/30/2023, score 895/1000

upvoted 2 times

  **mlantonis**  3 years, 7 months ago

Full image: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_062\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_062_Image_0001.jpg)

upvoted 79 times

  **mlantonis** 3 years, 7 months ago

Box 1: Service bus queue

You are developing a back-end Azure App Service that scales based on the number of messages contained in a Service Bus queue.

Box 2: ActiveMessage Count

ActiveMessageCount: Number of messages in the queue or subscription that are in the active state and ready for delivery.

Box 3: Average

For special metrics such as Storage or Service Bus Queue length metric, the threshold is the average number of messages available per current number of instances.

Box 4: Less than or equal to

You need to add a new rule that will continuously scale down the App Service, as long as the scale up condition is not met.

Box 5: Decrease count by

upvoted 38 times


  **mlantonis** 3 years, 7 months ago

Reference:

<https://docs.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-best-practices#considerations-for-scaling-threshold-values-for-special-metrics>

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/message-counters>



upvoted 9 times

  **warchoon** 1 year, 10 months ago

<https://learn.microsoft.com/en-us/azure/azure-monitor/autoscale/autoscale-best-practices#considerations-for-scaling-when-multiple-rules-are-configured-in-a-profile>

points Less than

upvoted 1 times

  **jay158** 3 years, 7 months ago

You are providing invaluable help, by clarifying the questions,

upvoted 12 times

  **neelkanths**  8 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

answer is correct...

The correct answers are

- 1) Service bus queue
  - 2) Active message count
  - 3) Average
  - 4) Less than or equal to
  - 5) Decrease count by
- upvoted 2 times

🗨️ 👤 **NightshadeRC** 1 year, 5 months ago

Had this question today: 2023-07-26

upvoted 3 times

🗨️ 👤 **AsuraX** 1 year, 11 months ago

Got this on my exam, January 10, 2023 (I passed)

My answer:

Box 1: Service bus queue

Box 2: Active Message Count

Box 3: Average

Box 4: Less than or equal to

Box 5: Decrease count by

upvoted 3 times

🗨️ 👤 **Lucifer14** 2 years ago

In 28-12-2022 exam, Answers.

Service Bus queue

active message count

count

less than equal to

decrease count

upvoted 1 times

🗨️ 👤 **at\_sayali** 2 years, 1 month ago

Got it on 10/2022, I chose as below:

Box 1: Service bus queue

Box 2: Active Message Count

Box 3: Average

Box 4: Less than or equal to

Box 5: Decrease count by

upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 3 months ago

Admin, Request to please update the full image. Current image is cut and options 4 and 5 are not visible

upvoted 4 times

🗨️ 👤 **meoukg** 2 years, 9 months ago

Got it on 03/2022, I chose as below:

Box 1: Service bus queue

Box 2: ActiveMessage Count

Box 3: Average

Box 4: Less than or equal to

Box 5: Decrease count by

upvoted 3 times

🗨️ 👤 **GigaCaster** 3 years, 5 months ago

Just a question but how are you going to reduce the count if there is now count to reduce?

upvoted 2 times

🗨️ 👤 **glam** 3 years, 7 months ago

The correct answers are

1) Service bus queue

2) Active message count

3) Average

4) Less than or equal to

5) Decrease count by

upvoted 4 times

🗨️ 👤 **Phenr** 3 years, 9 months ago

I believe the answers are:

1st box: as given, since we're analyzing the state of messages in the Service Bus Queue

2nd box: as given, because Active Message Count will show us just the count of messages that are in the active state and ready for delivery, while Message Count would show us all the messages.

3rd box: Average, since we should stay consistent with the Scale Up time grain, which is also "Average" - this way, when we use "average" on both, we have consistent and clear conditions when to go up and when to go down

upvoted 15 times

🗨️ 👤 **hobob** 3 years, 9 months ago

I think the 3rd option should be "Average" since the scale up rule is based off average, it makes sense to make the scale down rule work on the same principal.

upvoted 6 times

🗨️ 👤 **Mike\_St** 3 years, 9 months ago

The image is cut so i cant see the bottom 2 dropdowns but they sound alright in the answer explanation bellow.

Also the 2nd box should be Message Count.

Reason:

In the question it is said:

"when the average queue length of unprocessed AND valid queue messages"

If it was just valid Queue messages it would have been enough to set it on Active Message Count... but they want Valid AND Unprocessed messages... so i think it should be Message Count.

(i am not 100% sure but unprocessed message should not count as an active message)

upvoted 3 times

🗨️ 👤 **qerem** 3 years, 9 months ago

Messagecount is false because it also contains deadlettered messages among others. Unprocessed messages are active because they are "ready to delivery". So i think "Active Message Count" is correct

upvoted 3 times

DRAG DROP -

You have an application that uses Azure Blob storage.

You need to update the metadata of the blobs.

Which three methods should you use to develop the solution? To answer, move the appropriate methods from the list of methods to the answer area and arrange them in the correct order.

Select and Place:

### Methods

Metadata.Add

SetMetadataAsync

FetchAttributesAsync

UploadFileStream

SetPropertiesAsync

### Answer Area

### Methods

Suggested Answer:

FetchAttributesAsync

UploadFileStream

### Answer Area

Metadata.Add

SetMetadataAsync

SetPropertiesAsync

Metadata.Add example:

```
// Add metadata to the dictionary by calling the Add method
metadata.Add("docType", "textDocuments");
```

SetMetadataAsync example:

```
// Set the blob's metadata.
await blob.SetMetadataAsync(metadata);
// Set the blob's properties.
await blob.SetPropertiesAsync();
```

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-properties-metadata>

 **Phenr**  4 years, 3 months ago

Since we're talking about updating the metadata,

- first we need to fetch it, to populate blob's properties and metadata (we want to update it - without fetching we would just set the new metadata):

FetchAttributesAsync

- second, we need to manipulate the metadatas to update them and the best fitting is

Metadata.Add

- third, we have to persist our changes. We can use a method that initiates an asynchronous operation to update the blob's metadata, which is

SetMetadataAsync

upvoted 262 times

 **zzt** 3 years, 11 months ago

Based on MS documentation referenced it should be:

```
// Get the blob's properties and metadata.
```

```
BlobProperties properties = await blob.GetPropertiesAsync();
```

```
// Add metadata to the dictionary by calling the Add method
```

```
properties.metadata.Add("docType", "textDocuments");
```

```
// Add metadata to the dictionary by using key/value syntax
properties.metadata["category"] = "guidance";
```

```
// Set the blob's metadata.
```

```
await blob.SetMetadataAsync(properties.metadata);
upvoted 21 times
```

🗨️ **altafpatel1984** 3 years, 7 months ago

But there is no option for GetProperties !

upvoted 2 times

🗨️ **coffecold** 2 years, 8 months ago

Yes there is, in version 12. The exam topic references for version 11 (legacy)

upvoted 3 times

🗨️ **azurelearner666** 4 years ago

This is the best response (the "official" response is wrong)

upvoted 4 times

🗨️ **rdemontis** 4 years, 3 months ago

correct!

upvoted 2 times

🗨️ **TakumaK** 4 years, 1 month ago

nicely explained!

upvoted 3 times

🗨️ **Tom87** **Highly Voted** 4 years, 2 months ago

I suppose there are two correct answers, depending on the version of Azure.Storage.Blobs. For v11: FetchAttributesAsync, Metadata.Add, SetMetadataAsync.

For v12: GetPropertiesAsync, Metadata.Add, SetMetadataAsync.

Just look here, there are two tabs with source code. One for v11, one for v12:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-properties-metadata?tabs=dotnet#set-and-retrieve-metadata>

upvoted 57 times

🗨️ **vtomy** 4 years ago

Correct. But GetPropertiesAsync not mentioned in answer options. We can go with FetchAttributesAsync

upvoted 14 times

🗨️ **xRiot007** 3 years ago

MS updated their pages.

Now, to set metadata, you need to add it either with metadata.Add, or metadata[key] = value.

After, you just need to save it, by calling SetMetadataAsync.

No other operations are required as we do not care what metadata is already on the blob set.

upvoted 1 times

🗨️ **macobuzi** 1 year, 10 months ago

Yeah, But the question forces you to pick 3 options. So FetchAttributeAsync is probably the closest one.

upvoted 1 times

🗨️ **Vichu\_1607** **Most Recent** 7 months, 3 weeks ago

Correct Order:

c. FetchAttributesAsync

a. Metadata.Add

b. SetMetadataAsync

upvoted 1 times

🗨️ **alora01** 1 year, 3 months ago

Sure, here are three methods you could use to update the metadata of blobs in Azure Blob storage-

SetBlobMetadata

SetMetadataAsync

SetPropertiesAsync

For more details - <https://sysconverter.com/blog/import-pst-to-shared-mailbox-office-365/>

upvoted 2 times

- 🗨️ 👤 **Mattt** 8 months, 1 week ago  
irrelevant link  
upvoted 1 times
- 🗨️ 👤 **onlyforheros** 1 year, 3 months ago  
Got it in exam on 13.03.2024. Score: 910.  
upvoted 2 times
- 🗨️ 👤 **arunkum1** 1 year, 6 months ago  
Got it in the exam 14/12/23. Went with given answer. All questions are from ExamTopics. Case study - VanArsdel, Ltd (11 questions)  
upvoted 2 times
- 🗨️ 👤 **AndySmith** 1 year, 7 months ago  
On exam 3-Nov-2023. Went with most-voted answer - 932/1000.  
1) FetchAttributesAsync  
2) Metadata.Add  
3) SetMetadataAsync  
upvoted 2 times
- 🗨️ 👤 **ENGs** 1 year, 8 months ago  
On my exam 2023-10 before the Update of the Exam  
upvoted 1 times
- 🗨️ 👤 **RuffBoii** 1 year, 9 months ago  
Had this on my exam today.  
upvoted 1 times
- 🗨️ 👤 **dddddd111** 1 year, 9 months ago  
I got this same question. Provided answers are correct. (Note: I failed the exam 20/Sept/23. I only scored 644 and I felt bad. I think because many questions here in Examtopics are not correct or already outdated. I suggest following the most voted answers and not rely on Examtopics answers. At the beginning of the exam, you will be asked which programming languages you want to use. C#/Python. I chose C#. Also, I just want to add that some questions here in the actual exams, but the choices are written and formatted differently. Please be aware of that. Goodluck. I feel bad for failing it, but I want to retake next month. I will try Python. T\_T  
upvoted 3 times
- 🗨️ 👤 **TechyNetty** 1 year, 9 months ago  
Sorry to hear that. I have exam soon. Were all the questions from ExamTopics> Any tips you would like to share please, Thanks  
upvoted 1 times
- 🗨️ 👤 **MedELKARMI** 1 year, 2 months ago  
Do the questions you found in the exam all come from here?  
upvoted 1 times
- 🗨️ 👤 **dddddd111** 1 year, 7 months ago  
I would say 30% percent of the questions here really shown to my exam. But after I purchased contributor access, I would say 90 percent of the questions here. In my first attempt, I almost lose my hope of passing because of the case study. So, I highly recommend you review case studies. really difficult. And like I said, some question and answer format are quite different from the actual exam. Good luck. I'm gonna retake before end of Nov.2023 because of the recent update from MS. AZ-204 was updated.  
upvoted 1 times
- 🗨️ 👤 **longnguyendh** 1 year, 11 months ago  
It should be FetchAttributesAsync, Metadata.Add, SetMetadataAsync  
upvoted 2 times
- 🗨️ 👤 **juanckar** 1 year, 11 months ago  
This was on the exam (July 2023). Went with fetch/metadata.add/setMetadata. Scored 917  
upvoted 4 times
- 🗨️ 👤 **DonH** 2 years ago  
Just for information: I just had this question on my AZ204 exam - 16-jun-2023.  
I barely made it (with only 767 points) so I can't inform anyone if this answer is correct or not, just stating that this is an actual exam question.  
upvoted 4 times
- 🗨️ 👤 **saravanasanthosh** 2 years, 6 months ago  
Got this in exam 12/30/2022  
upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 6 months ago  
thanks to all who have mentioned the exam dates  
upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago  
FetchAttributesAsync  
Metadata.Add  
SetMetadataAsync  
upvoted 3 times

🗨️ 👤 **elequiel** 2 years, 8 months ago  
Got it in exam 20/10/2022  
upvoted 2 times



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 are developing an Azure solution to collect point-of-sale (POS) device data from 2,000 stores located throughout the world. A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data. You must store the device data in Azure Blob storage. Device data must be correlated based on a device identifier. Additional stores are expected to open in the future.

You need to implement a solution to receive the device data.

Solution: Provision an Azure Event Grid. Configure the machine identifier as the partition key and enable capture.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: A**

Reference:

<https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services>

Community vote distribution


 B (100%)

  **bbijls**  4 years, 3 months ago

I don't know who gave all the answer, but that person definitely failed the exam..  
upvoted 149 times

  **kondapaturi** 4 years ago



what is the answer?  
upvoted 1 times

  **Anil4Az** 4 years, 1 month ago


Good one... :)  
upvoted 4 times

  **AnuSubramaniam** 3 years, 10 months ago

Its Event Hub not, Grid.  
upvoted 25 times

  **basquiat** 1 year, 10 months ago

I honestly think that this is on purpose so people actually discuss the questions. Reverse psychology. LOL  
upvoted 6 times

  **Occ50bf** 10 months, 3 weeks ago



If they posted the real answers they would have Microsoft on their tail immediately.  
upvoted 1 times

  **Dinima**  4 years, 3 months ago

I think event hub. Azure Event Hub can be used to get the messages from the various devices. Azure Event Hub capture can then be used to persist the events to Azure Blob storage.  
upvoted 53 times

  **Tom87** 4 years, 2 months ago

I agree. The solution says "enable capture". I have found capture only for Event Hub, not for Event Grid.  
<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-capture-overview>  
upvoted 12 times

  **Esward** 2 years, 5 months ago

Agreed  
upvoted 1 times

🗨️ **Vichu\_1607** Most Recent 7 months, 3 weeks ago

No, the solution does not meet the goal.

Azure Event Grid is primarily used for event-driven architectures and is designed to handle events and notifications. It is not suitable for directly ingesting and storing large amounts of data, such as the 2 MB of data per device per day from 2,000 stores.

upvoted 1 times

🗨️ **zeugo** 1 year, 6 months ago

Its Eventhub, grid is primarily for event routing, allowing other resources like functions to react to changes from a blobstorage etc. Event hub is the choice for streaming telemetry data from thousands of sources to a blobstorage. Here is a digestable documentation for understanding the difference between Hub, Grid and ServiceBus: <https://arindam-das.medium.com/demystifying-azures-eventing-services-a-comparison-of-event-hub-event-grid-and-service-bus-d578693dcf16>

upvoted 2 times

🗨️ **Samuel316** 1 year, 8 months ago

It should be EventHub, not EventGrid.

upvoted 1 times

🗨️ **ENGs** 1 year, 8 months ago

Had this series of questions On my exam 2023-10 before the Update of the Exam

upvoted 1 times

🗨️ **macobuzi** 1 year, 10 months ago

The answer is Azure Event Hub, not Event Grid.

upvoted 5 times

🗨️ **KarthikCS** 2 years, 2 months ago

How come almost every answer marked by examtopic is incorrect. Whats the purpose of marking answers then if you have to follow discussion thread for each question?

upvoted 9 times

🗨️ **dddddd111** 1 year, 7 months ago

Yeah, it is confusing. I failed in my first attempt. scored 644 only.

upvoted 2 times

🗨️ **Kluk\_Kluk** 2 years, 3 months ago

"device data from 2,000 stores located throughout the world." It is a distributed data streaming. Answer is EventHub

upvoted 5 times

🗨️ **fkaracan** 2 years, 4 months ago

Selected Answer: B

Given solution does not meet the goal because you are using Event Grid instead of Event Hubs. Event Grid does not have a capture feature and it does not store data in Blob storage. You need to use Event Hubs as your source and enable capture with Blob storage as your destination.

upvoted 4 times

🗨️ **alexein74** 2 years, 5 months ago

Selected Answer: B

B. No.

Azure Event Grid is an event routing service that allows you to handle events from various Azure services and your own applications. It can be used to send events from an application to multiple subscribers, but it is not well suited for receiving data from thousands of devices and storing them in Azure Blob storage.

Azure Event Grid is also not meant for long-term data storage and it is not a good fit for this scenario, where large amounts of data need to be stored and correlated based on a device identifier.

upvoted 4 times

🗨️ **cemergen** 2 years, 6 months ago

Selected Answer: B

B is the correct answer

upvoted 1 times

🗨️ **Knightie** 2 years, 10 months ago

Selected Answer: B

The reason for No is due to the max no. of partition. The question demands  $5 \times 2000 = 10000$  partition, which is more than all tiers available.

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-quotas#basic-vs-standard-vs-premium-vs-dedicated-tiers>



upvoted 3 times

  **Satish\_Babu** 2 years, 10 months ago

**Selected Answer: B**

EventHub not EventGrid

upvoted 1 times

  **sahirl** 2 years, 11 months ago

**Selected Answer: B**

B is the correct answer



upvoted 1 times

  **AZAdam22** 2 years, 11 months ago

**Selected Answer: B**

No - Capture is a feature of EventHub, not EventGrid.

upvoted 2 times

  **Dani\_ac7** 2 years, 12 months ago

**Selected Answer: B**

is event hub, not grid

upvoted 2 times

You develop Azure solutions.

A .NET application needs to receive a message each time an Azure virtual machine finishes processing data. The messages must NOT persist after being processed by the receiving application.

You need to implement the .NET object that will receive the messages.

Which object should you use?

- A. QueueClient
- B. SubscriptionClient
- C. TopicClient
- D. CloudQueueClient

**Suggested Answer: D**

A queue allows processing of a message by a single consumer. Need a CloudQueueClient to access the Azure VM.


Incorrect Answers:

B, C: In contrast to queues, topics and subscriptions provide a one-to-many form of communication in a publish and subscribe pattern. It's useful for scaling to large numbers of recipients.

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-queues-topics-subscriptions>

Community vote distribution



**mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: A

Azure.Storage.Queues.QueueClient: .NET v12

Azure.Storage.Queues.CloudQueueClient: .NET v11 (Legacy)

So, the question is really about what kind of queue message tool you should use. And the key word here is that "message must NOT persist after being processed".

Azure.Storage.Queues.QueueClient supports "At-Most-Once" deliver mode, while Azure.Storage.Queues.CloudQueueClient doesn't.

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/azure.storage.queues.queueclient?view=azure-dotnet>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueueclient?view=azure-dotnet-legacy>

upvoted 93 times

**Ody** 5 months, 3 weeks ago

Copilot says: Neither `Azure.Storage.Queues.QueueClient` nor `Azure.Storage.Queues.CloudQueueClient` supports the **At-Most-Once** delivery mode.

Both of these clients operate on an **At-Least-Once** delivery model, meaning a message might be delivered more than once if it is not explicitly deleted after being processed.

If you need **At-Most-Once** delivery, Azure Service Bus with the **ReceiveAndDelete** mode is a better option

upvoted 1 times

**guchao2000** 2 years, 5 months ago

It's QueueClient of Service Bus.

upvoted 4 times

🗳️ 👤 **MiraA** 3 years, 8 months ago

It seems the CloudQueueClient is a legacy class and is used for creation of QueueClient instances.

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.storage.queue.cloudqueueclient.getqueuereference>

upvoted 2 times

🗳️ 👤 **SSR999** 1 year, 5 months ago

but when you do receive message on storage queue it wont delete message automatically

upvoted 1 times

🗳️ 👤 **Spooky7** Highly Voted 4 years, 1 month ago

First of all - question is not precise as we don't know which QueueClient they are asking about. There are two options:

- Microsoft.AzureServiceBus.QueueClient?

- Azure.Storage.Queues.QueueClient?

I would say it is about Microsoft.AzureServiceBus.QueueClient as the difference between Azure.Storage.Queues.CloudQueueClient (v12) is just a legacy version of the Azure.Storage.Queues.QueueClient (v11)

So the question is really about what kind of queue message tool you should use. And the key word here is that "message must NOT persist after being processed". So correct answer would be Microsoft.AzureServiceBus.QueueClient (A) as it supports "At-Most-Once" deliver mode while Azure.Storage.Queues.CloudQueueClient doesn't.

upvoted 19 times

🗳️ 👤 **basquiat** 1 year, 10 months ago

I agree. Plus that Azure.Storage.Queues.QueueClient AFAIK is more appropriate for data intensive scenarios.

upvoted 1 times

🗳️ 👤 **wafa\_chaari** Most Recent 5 months ago

Selected Answer: A

QueueClient allows us to manipulate a queue and its msgs. we can add, update delete msgs in a queue. we can also receive msgs from a queue which can then be processed and deleted.

upvoted 1 times

🗳️ 👤 **Vichu\_1607** 7 months, 3 weeks ago

Selected Answer: A

QueueClient is the answer

upvoted 1 times

🗳️ 👤 **8ac3742** 10 months, 4 weeks ago

The answer is "QueueClient".The CloudQueueClient from Azure Storage Queue requires manual deletion of messages, while the QueueClient from Azure Service Bus can use the ReceiveAndDelete mode to automatically delete messages after processing.

upvoted 1 times

🗳️ 👤 **ShoaibAnwar** 1 year, 3 months ago

Selected Answer: B

SubscriptionClient Given the requirement for messages to not persist after being processed and the implied need for a publish/subscribe model

upvoted 2 times

🗳️ 👤 **raymond\_abcd** 1 year, 5 months ago

Selected Answer: A

Think it the servicebus namespace because it is messaging not the storagequeue namespace. The QueueClient is deprecated for Servicebus. It is not in the answers but it should be ServiceBusClient. So in this case I would go for the deprecated object answer A

upvoted 2 times

🗳️ 👤 **RikinPatel** 1 year, 9 months ago

Selected Answer: A

@Admin Please correct answer.

As per MS Doc and ChatGPT correct ans is A so please correct ans and its explanation

upvoted 2 times

🗳️ 👤 **alexein74** 2 years, 5 months ago

Selected Answer: A

A. QueueClient

upvoted 1 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

Selected Answer: A



A. QueueClient

upvoted 1 times

  **gmishra88** 2 years, 8 months ago

Microsoft creates all these random naming convention showing how disorganized they are in individual islands and they dare to ask questions along this fault lines.



upvoted 7 times

  **Perplex** 2 years, 11 months ago

**Selected Answer: A**

A is correct, the new API calls it QueueClient. See also most upvoted answer.

upvoted 3 times

  **ns4098** 9 months, 2 weeks ago

QueueClient: This is used for Azure Service Bus Queues, which follows a one-to-one messaging pattern, meaning only one receiver can process a message. However, messages in a queue persist until they are explicitly deleted or expire, which does not fit the requirement.

upvoted 1 times

  **xRiot007** 3 years ago

A - version 12

D - "legacy" version 11



upvoted 1 times

  **Bartimaeus** 3 years, 4 months ago

What about SeviceBus QueueClient in the old API?

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.servicebus.queueclient?view=azure-dotnet>

upvoted 1 times

  **idroj** 3 years, 4 months ago

**Selected Answer: A**

A is the correct answer

upvoted 3 times

  **Ccastan1** 3 years, 4 months ago

So the explanation for why A is wrong is incorrect? It states you cannot access the VM using A

upvoted 1 times

  **prabhjot** 4 years, 1 month ago

yes A is correct

upvoted 1 times

## DRAG DROP -

You are maintaining an existing application that uses an Azure Blob GPv1 Premium storage account. Data older than three months is rarely used. Data newer than three months must be available immediately. Data older than a year must be saved but does not need to be available immediately.

You need to configure the account to support a lifecycle management rule that moves blob data to archive storage for data not modified in the last year.

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

## Answer Area

Upgrade the storage account to GPv2

Create a new GPv2 Standard account and set its default access tier level to cool

Change the storage account access tier from hot to cool

Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account



**Suggested Answer:**

Actions	Answer Area
	Upgrade the storage account to GPv2
Create a new GPv2 Standard account and set its default access tier level to cool	Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account
	Change the storage account access tier from hot to cool

Step 1: Upgrade the storage account to GPv2

Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering.

You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

Step 3: Change the storage account access tier from hot to cool

Note: Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

**sien** Highly Voted 4 years, 2 months ago

Since we already have a premium P1 account with gpv1. Why not:

- Upgrade the existing one to GPv2
- Create a new GPV2 standard account with default access level to cool

- And then copy archive data to the GPV2 and delete the data from original storage account.

That makes sense to me.

upvoted 124 times

🗨️ 👤 **jay158** 4 years ago

Is there any requirement in question, which says set default access tier to COOL?

Please clarify.

upvoted 2 times

🗨️ 👤 **lighting** 4 years ago

I have 2 ideas:

1.

- One HOT for newer than 3 months

- One COOL for older than 3 months and Archive data.

2.

- One HOT for non-archived data. Can be accessed immediately. Because they just said "Older than 3 months data are rarely accessed" but didn't tell us anything about can it be accessed immediately.

- One COOL for archived data. Archived tier can just be set at blob level. "Only the hot and cool access tiers can be set at the account level. The archive access tier can only be set at the blob level" <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

upvoted 4 times

🗨️ 👤 **WillPassExam** 4 years, 2 months ago

does this mean at the end, we have 2 GPv2 storage accounts, one access level is cool (for archive data) and the other is hot?

upvoted 2 times

🗨️ 👤 **sien** 4 years, 2 months ago

I would say so.

upvoted 1 times

🗨️ 👤 **sien** 4 years, 2 months ago

Also look at this blog:

<https://www.apptio.com/blog/essential-guide-azure-blob-storage-pricing/>

Only GPv2 and Blob storage accounts support tiering. If you are using GPv1, and you want to leverage tiering, convert your account to GPv2 through the Azure portal.

upvoted 4 times

🗨️ 👤 **ferut** 4 years, 1 month ago

My choice as well

upvoted 1 times

🗨️ 👤 **Chked** 3 years, 11 months ago

I agree. But can the first step be the last?

- Create a new GPV2 standard account with default access level to cool

- And then copy archive data to the GPV2 and delete the data from original storage account.

- Upgrade the existing one to GPv2

upvoted 1 times

🗨️ 👤 **mlantonis** Highly Voted 4 years, 1 month ago

Step 1: Upgrade the storage account to GPv2

Object storage data tiering between hot, cool, and archive is supported in Blob Storage and General Purpose v2 (GPv2) accounts. General Purpose v1 (GPv1) accounts don't support tiering. You can easily convert your existing GPv1 or Blob Storage accounts to GPv2 accounts through the Azure portal.

Step 2: Create a new GPV2 standard account with default access level to cool

Step 3: Copy the data to be archived to a Standard GPv2 storage account and then delete the data from the original storage account

upvoted 23 times

🗨️ 👤 **tosm** 3 years ago

Agreed: 1. upgrade storage account 1 from GPv1 -> v2 (hot), 2. create storage account 2 GPv2 (cool) 3. copy data from account 1 to account 2.

End result: account 1 (hot), account 2(cool and archive) both accounts can set lifecycle policy

upvoted 4 times

🗨️ 👤 **mlantonis** 4 years, 1 month ago



Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>


<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-upgrade?tabs=azure-portal>

upvoted 2 times

  **mlantonis** 4 years, 1 month ago

Although Step 3 seems unusual and not necessary.

upvoted 4 times

  **ensa** 3 years, 8 months ago

every time with good explanation but step 3 needed because why not transfer the old data that needed to new one and delete the old one for saving cost

upvoted 1 times

  **8ac3742**  10 months, 4 weeks ago

In real practice, I just keep one account: upgrade the account to GPV2 which supports tier, change the account tier to cool since the data in the account is infrequently used, and achieve the blob which is not modified for more than 1 year.

upvoted 1 times

  **Christian\_garcia\_martin** 10 months, 4 weeks ago



Upgrade , Create and Copy

upvoted 1 times

  **11\_NickName\_11** 1 year, 7 months ago

Why to upgrade the storage account to GPv2?

upvoted 2 times

  **AlbertoBT** 1 year, 7 months ago

The requirement is just: You need to configure the account to SUPPORT a lifecycle management rule.

So you only need "Upgrade the existing one to GPv2"

No more steps from the list are required

Then you could configure the lifecycle management rule.

Once you apply the rule, the files tier will be changed automatically

upvoted 2 times

  **narenazure** 2 years, 2 months ago

Azure Blob storage lifecycle management offers a rich, rule-based policy for General Purpose v2 and Blob storage accounts.

upvoted 1 times

  **deepak\_26** 2 years, 3 months ago

Create , change , copy



Data stored in a premium block blob storage account cannot be tiered to Hot, Cool, or Archive using Set Blob Tier or using Azure Blob Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the Hot tier in a different account using the Put Block From URL API or a version of AzCopy that supports this API.

upvoted 3 times

  **motekim** 2 years, 3 months ago



Just for reference: GPv2 supports Hot, Cool, and Archive tiers

upvoted 2 times

  **dimsok** 2 years, 4 months ago

Upgrade-Create-Copy

upvoted 2 times

  **JamieS** 2 years, 5 months ago

Seems like we don't have the entire world view, and that there's already an existing standard v2 account elsewhere. If that's the case we're starting with v1 storage where the data currently is and we also have a default v2 storage that exists.

We upgrade the current v1 to v2 so we can access lifetime management

we copy the data to be archived to the standard v2 (that already existed, and by default has hot tier)

then we set the tier to be cool (where the archive data was just copied to).

I think the current answer is correct.

upvoted 2 times

🗨️ 👤 **micro9000** 2 years, 7 months ago

The answer doesn't make any sense to me.

in the 2nd step, it says copy the data from old storage account to new one, it means there are 2 storage accounts, why you need to upgrade the existing?

upvoted 2 times

🗨️ 👤 **Akiu** 2 years, 8 months ago

Upgrade to GPv2

Copy data

Change tier to cool

You can't create the second account in a cool tier because of this:

Data stored in a premium block blob storage account cannot be tiered to hot, cool, or archive using Set Blob Tier or using Azure Blob Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the hot tier in a different account using ...

<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#blob-lifecycle-management>

upvoted 1 times

🗨️ 👤 **Azziet** 3 years ago

Upgrade

Create

Copy

upvoted 7 times

🗨️ 👤 **Evo\_Morales** 3 years, 2 months ago

Agree with voting/answer, but the question itself seems flawed. Need to learn not to read anything else into the questions and not add steps/requirements that yes, would make sense.

upvoted 3 times

🗨️ 👤 **petitbilly** 3 years, 3 months ago

Got it in exam 03/22

upvoted 5 times

🗨️ 👤 **massnonn** 3 years, 4 months ago

for me not have sense - Upgrade the existing one to GPv2

So:

-create ..

-change..

-copy

upvoted 1 times

You develop Azure solutions.

You must connect to a No-SQL globally-distributed database by using the .NET API.

You need to create an object to configure and execute requests in the database.

Which code segment should you use?

- A. `new Container(EndpointUri, PrimaryKey);`
- B. `new Database(EndpointUri, PrimaryKey);`
- C. `new CosmosClient(EndpointUri, PrimaryKey);`

**Suggested Answer: C**

Example:

```
// Create a new instance of the Cosmos Client
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
//ADD THIS PART TO YOUR CODE
await this.CreateDatabaseAsync();
```

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql-api-get-started>

Community vote distribution

C (100%)

 **Mr2302682** Highly Voted 4 years, 10 months ago

The Answer is correct!

upvoted 49 times

 **mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: C

Azure Cosmos DB is a fully managed NoSQL database for modern app development. Single-digit millisecond response times, and automatic and instant scalability, guarantee speed at any scale.

```
// Create a new instance of the Cosmos Client
this.cosmosClient = new CosmosClient(EndpointUri, PrimaryKey)
//ADD THIS PART TO YOUR CODE
await this.CreateDatabaseAsync();
```

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient?view=azure-dotnet>

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql-api-get-started>


upvoted 23 times

 **wafa\_chaari** Most Recent 5 months ago

**Selected Answer: C**


cosmosclient is the entry point for interacting with Azure Cosmos DB

upvoted 1 times

 **Tej\_007** 9 months, 3 weeks ago

wow. finally one correct answer.

upvoted 1 times

 **dimsok** 2 years, 4 months ago

I dont get this question, Cosmos DB is a No-SQL globally-distributed database, but there are others that .NET can connect.

upvoted 2 times

- 🗳️ 👤 **Jorax** 2 months, 2 weeks ago  
Yes, but "You develop Azure solutions."  
upvoted 1 times
- 🗳️ 👤 **Nhiendo** 2 years, 5 months ago  
Correct Answer: C  
upvoted 1 times
- 🗳️ 👤 **AvinashVarma** 2 years, 6 months ago  
**Selected Answer: C**  
CosmosClient(String, String, CosmosClientOptions)  
  
<https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient?view=azure-dotnet>  
upvoted 1 times
- 🗳️ 👤 **OPT\_001122** 2 years, 7 months ago  
**Selected Answer: C**  
C. new CosmosClient(EndpointUri, PrimaryKey);  
upvoted 1 times
- 🗳️ 👤 **rohitpatil113** 2 years, 10 months ago  
**Selected Answer: C**  
The Answer is correct!  
upvoted 1 times
- 🗳️ 👤 **Eltooth** 3 years ago  
**Selected Answer: C**  
C is correct answer.  
Global fully managed No-SQL DB = CosmosDB  
upvoted 2 times
- 🗳️ 👤 **petitbilly** 3 years, 3 months ago  
Got it in exam 03/22  
upvoted 2 times
- 🗳️ 👤 **Freidrich** 3 years, 4 months ago  
**Selected Answer: C**  
The correct answer is C.  
upvoted 1 times
- 🗳️ 👤 **oescm** 3 years, 4 months ago  
Got this one 02/2022. Went with highly voted answer.  
upvoted 5 times
- 🗳️ 👤 **Mev4953** 3 years, 5 months ago  
Got this in the exam 01/22  
upvoted 5 times
- 🗳️ 👤 **lugospod** 3 years, 5 months ago  
Got this one 01/2022. Went with most voted (to avoid writing answers again)  
upvoted 5 times
- 🗳️ 👤 **sanjayrawat** 3 years, 8 months ago  
Answer : C  
upvoted 2 times
- 🗳️ 👤 **debanjan10** 3 years, 8 months ago  
Order:  
Cosmos Client > Database > Container > Item  
upvoted 3 times

You have an existing Azure storage account that stores large volumes of data across multiple containers.

You need to copy all data from the existing storage account to a new storage account. The copy process must meet the following requirements:

- ⇒ Automate data movement.
- ⇒ Minimize user input required to perform the operation.
- ⇒ Ensure that the data movement process is recoverable.

What should you use?

- A. AzCopy
- B. Azure Storage Explorer
- C. Azure portal
- D. .NET Storage Client Library

**Suggested Answer: A**

You can copy blobs, directories, and containers between storage accounts by using the AzCopy v10 command-line utility.

The copy operation is synchronous so when the command returns, that indicates that all files have been copied.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy>

Community vote distribution

A (100%)

🗳️ **kondapaturi** Highly Voted 3 years, 6 months ago

Answer – AzCopy, The Azcopy tool can be used to copy data from one storage account to another. You can use the tool within automation scripts to ensure the data can be copied automatically.

upvoted 42 times

🗳️ **mlantonis** Highly Voted 3 years, 7 months ago

Correct Answer: C

Azure Storage Explorer uses AzCopy to perform all of its data transfer operations. But in this questions, there is a requirement to minimize user interaction which is why AzCopy is more appropriate.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-blobs-copy>

upvoted 10 times

🗳️ **alperc** 3 years, 1 month ago

please stop commenting..you are always misanswering.

upvoted 8 times

🗳️ **altafpatel1984** 3 years, 1 month ago

But read very first requirement. It needs to be Automated. I agree with alperc, you are misanswering, even such a silly question.

upvoted 5 times

🗳️ **azurelearner666** 3 years, 6 months ago

AzCopy is A... please do not write confusing responses...

response is already correct, A - AzCopy

upvoted 47 times

🗳️ **jobolesoniha1** Most Recent 9 months, 1 week ago

**Selected Answer: A**

It's no brainer. AzCopy is the answer with no or minimal human intervention.

upvoted 1 times

🗄️ 👤 **NombreFalso** 1 year, 10 months ago

**Selected Answer: A**

Typical AzCopy  
upvoted 2 times

🗄️ 👤 **AvinashVarma** 2 years ago

**Selected Answer: A**

AzCopy  
Reference: <https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-v10>  
upvoted 2 times

🗄️ 👤 **OPT\_001122** 2 years, 1 month ago

**Selected Answer: A**

A. AzCopy  
upvoted 1 times

🗄️ 👤 **Unpredictable\_m** 2 years, 2 months ago

Correct answer is A  
AzCopy is a command-line utility that you can use to copy blobs or files to or from a storage account  
upvoted 1 times

🗄️ 👤 **Chaitya** 2 years, 4 months ago

AzCopy  
upvoted 1 times

🗄️ 👤 **ghuna** 2 years, 4 months ago

**Selected Answer: A**

AzCopy is best option to move data around in blob storage  
upvoted 1 times

🗄️ 👤 **xRiot007** 2 years, 6 months ago

You will use AzCopy. It can be automated in a script, reducing user interaction and you can also recover  
upvoted 1 times

🗄️ 👤 **Eltooth** 2 years, 6 months ago

**Selected Answer: A**

A is the correct answer.  
Everytime.  
upvoted 1 times

🗄️ 👤 **meoukg** 2 years, 9 months ago

Got it on 03/2022, I chose the A. AzCopy  
upvoted 2 times

🗄️ 👤 **idroj** 2 years, 10 months ago

**Selected Answer: A**

AzCopy is the correct answer  
upvoted 3 times

🗄️ 👤 **Exocsinstr** 3 years ago

<https://docs.microsoft.com/en-us/azure/storage/common/storage-use-azcopy-configure>  
Ans: should be AzCopy. This link show you can recover the AzCopy from where it failed.

`azcopy jobs resume <job-id> --source-sas="<sas-token>" --destination-sas="<sas-token>"`  
Azure Data Factory in the Azure portal may need much configuration by user to setup the pipeline to perform copy.  
upvoted 1 times

🗄️ 👤 **francis6170** 3 years, 6 months ago

got this in the exam :)  
upvoted 6 times

🗄️ 👤 **debanjan10** 3 years, 2 months ago

A lot of these questions are common in the exam it seems.  
upvoted 1 times

🗄️ 👤 **jeet1985** 3 years, 6 months ago



Then ? please suggest the correct answer?

upvoted 4 times

  **UnknowMan** 3 years, 7 months ago

The Answer is correct!

upvoted 4 times

  **glam** 3 years, 7 months ago

correct

upvoted 3 times

## DRAG DROP -

You are developing a web service that will run on Azure virtual machines that use Azure Storage. You configure all virtual machines to use managed identities.

You have the following requirements:

- ⇒ Secret-based authentication mechanisms are not permitted for accessing an Azure Storage account.
- ⇒ Must use only Azure Instance Metadata Service endpoints.

You need to write code to retrieve an access token to access Azure Storage. To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

## Code segment 1

`http://localhost:50342/oauth2/token`

`http://169.254.169.254:50432/oauth2/token`

`http://localhost/metadata/identity/oauth2/token`

`http://169.254.169.254/metadata/identity/oauth2/token`

## Code segment 2

`XDocument.Parse(payload);`

`new MultipartContent(payload);`

`new NetworkCredential("Azure", payload);`

`JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);`

## Answer Area

```
var url = " Code segment 1 ";
var queryString = "...";
var client = new HttpClient();
var response = await client.GetAsync(url + queryString);
var payload = await response.Content.ReadAsStringAsync();

return Code segment 2
```

## Suggested Answer:

## Code segment 1

`http://localhost:50342/oauth2/token`

`http://169.254.169.254:50432/oauth2/token`

`http://localhost/metadata/identity/oauth2/token`

## Code segment 2

`XDocument.Parse(payload);`

`new MultipartContent(payload);`

`new NetworkCredential("Azure", payload);`

## Answer Area

```
var url = " http://169.254.169.254/metadata/identity/oauth2/token ";
var queryString = "...";
var client = new HttpClient();
var response = await client.GetAsync(url + queryString);
var payload = await response.Content.ReadAsStringAsync();

return JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);
```

Azure Instance Metadata Service endpoints "/oauth2/token"

Box 1: `http://169.254.169.254/metadata/identity/oauth2/token`

Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended):

GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' HTTP/1.1

Metadata: true

Box 2: `JsonConvert.DeserializeObject<Dictionary<string, string>>(payload);`

Deserialized token response; returning access code.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token>

<https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code>



Box 1: `http://169.254.169.254/metadata/identity/oauth2/token`

Sample request using the Azure Instance Metadata Service (IMDS) endpoint (recommended):

GET 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' HTTP/1.1

Metadata: true

Box 2: `JsonConvert.DeserializeObject<Dictionary<string,string>>(payload);`



Deserialized token response; returning access code.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token>

<https://docs.microsoft.com/en-us/azure/service-fabric/how-to-managed-identity-service-fabric-app-code>

upvoted 70 times

  **Knightie** 2 years, 10 months ago

IMDS is a REST API that's available at a well-known, non-routable IP address ( 169.254.169.254 ). You can only access it from within the VM.

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/instance-metadata-service?tabs=windows>

upvoted 7 times

  **warchoon** 2 years, 4 months ago

Powerful Words And The Magic Of Abracadabra :)

upvoted 2 times



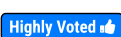
  **130nk3r5** 1 year, 6 months ago

Got this today.

Went with answer here.

Score 927

upvoted 3 times

  **dddddd111**  1 year, 9 months ago

I got this same question. Provided answers are correct. (Note: I failed the exam 20/9/23. I only scored 644 and I felt bad. I think because many questions here in Examtopics are not accurate. I suggest following the most voted answers and don't just not rely on Examtopics answers. At the beginning of the exam, you will be asked which programming languages you want to use. C#/Python. I chose C#. Also, I just want to add that some questions here are really in the actual exams, but the choices are written and formatted differently. Please be aware of that. Goodluck. I feel bad for failing it, but I want to retake next month. I will try Python. T\_T



upvoted 18 times

  **MasonZ** 10 months, 3 weeks ago

Man I respect you for the fact that you are the only one commenting, with honest, that you failed the exam, obviously all other people (or bots) got 925 points and passed the exam.

All the best and wish you passed the second attempt :D

upvoted 3 times

  **spahno** 1 year, 3 months ago

Hey man, heads up! 644 is a decent score I'd say. Did you pass the test the 2nd time? :)



upvoted 1 times

  **Kanasan**  1 year, 10 months ago

Correct!

<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token>

upvoted 2 times

  **mihailos** 1 year, 10 months ago



Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 1 times

  **macobuzi** 1 year, 10 months ago

Answer is correct!

upvoted 1 times

  **juanckar** 1 year, 11 months ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 3 times

🗨️ 👤 **juanckar** 1 year, 11 months ago

This was on the exam (July 2023). Went with proposed. Scored 917

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 4 times

🗨️ 👤 **Esward** 2 years, 4 months ago

It was there in 13 Feb 2023 exam

upvoted 6 times

🗨️ 👤 **martijn44** 2 years, 8 months ago

I was taught to use \$IDENTITY\_ENDPOINT but indeed that always seems to refer to the 169.254.169.254 address

upvoted 1 times

🗨️ 👤 **mijkol** 3 years ago

I get confused because of the IP, I get we shouldn't go for the local one, but why is this an AIPIPA IP adress?

is there a reason for that?

upvoted 2 times

🗨️ 👤 **Rini100** 3 years, 2 months ago

Got this on 20 Apr 2022

upvoted 2 times

🗨️ 👤 **kondapaturi** 4 years ago

BOX1 - <http://169.254.169.254/metadata/identity/oauth2/token> ,To get the metadata from the local service on the machine, the right URL is <http://169.254.169.254/metadata/identity/oauth2/token>

BOX2 – JsonConvert.DeserializeObject<string,string>(payload)

You can deserialize the response using the JsonConvert.DeserializeObject method. You can then get a dictionary collection and then get the access key from there.

upvoted 5 times

🗨️ 👤 **goatlord** 4 years ago

Correct

upvoted 2 times

🗨️ 👤 **UnknowMan** 4 years, 1 month ago

Correct ! => <https://docs.microsoft.com/fr-fr/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token>

upvoted 2 times

🗨️ 👤 **glam** 4 years, 1 month ago

correct

upvoted 2 times

🗨️ 👤 **clarionprogrammer** 4 years, 2 months ago

It is correct.

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/how-to-use-vm-token>

upvoted 12 times

## DRAG DROP -

You are developing a new page for a website that uses Azure Cosmos DB for data storage. The feature uses documents that have the following format:

```
{
  "name": "John",
  "city": "Seattle"
}
```

You must display data for the new page in a specific order. You create the following query for the page:

```
SELECT*
FROM People p
ORDER BY p.name, p.city DESC
```

You need to configure a Cosmos DB policy to support the query.

How should you configure the policy? To answer, drag the appropriate JSON segments to the correct locations. Each JSON segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

**JSON segments****Answer Area**

```
{
  "automatic": true,
  "ngMode": "Consistent",
  "includedPaths": [
    {
      "path": "/*"
    }
  ], "excludedPaths": [],
  "": [
    {
      "path": "/name", "order": "descending"
    },
    {
      "path": "/city", "order": " "
    }
  ]
}
```

#### Suggested Answer:

##### JSON segments

orderBy

sortOrder

ascending

descending

compositeIndexes

##### Answer Area

```
{
  "automatic": true,
  "ngMode": "Consistent",
  "includedPaths": [
    {
      "path": "/*"
    }
  ], "excludedPaths": [],
  "compositeIndexes": [
    {
      "path": "/name", "order": "descending"
    },
    {
      "path": "/city", "order": "descending"
    }
  ]
}
```

Box 1: compositeIndexes -

You can order by multiple properties. A query that orders by multiple properties requires a composite index.

Box 2: descending -

Example: Composite index defined for (name ASC, age ASC):

It is optional to specify the order. If not specified, the order is ascending.


```
{
  "automatic":true,
  "indexingMode":"Consistent",
  "includedPaths":[
    {
      "path": "/*"
    }
  ],
  "excludedPaths":[],
  "compositeIndexes":[
    {
      "path": "/name",
    },
    {
      "path": "/age",
    }
  ]
}
```

 **kayleena93**  4 years, 7 months ago

ORDER BY queries on multiple properties:

The composite index also supports an ORDER BY clause with the opposite order on all paths.



So I think it's about reversed index to the query. Answer should be 'ascending'. You cannot support ASC (default), DESC query with DESC, DESC index.  
upvoted 92 times

 **john4p** 3 years, 6 months ago

The problem here is the SQL that makes many people think that  
"ORDER BY p.name, p.city DESC"  
means it's ordered by name and city both descending.  
But the DESC only applies to city. name is ASC - this would be less confusing:  
"ORDER BY p.name ASC, p.city DESC"



Thus in the JSON you can only state ascending+descending or the opposite: descending+ascending.  
Since descending for name is already set the answer is "ascending".



At first I had misread the SQL wrong myself and didn't understand kayleena's comment right away.  
upvoted 37 times




  **Dobby92** 1 year, 5 months ago  
I just tried it.



- 1) Using Indexing Policy (name -> descending, city -> descending) gives an error:  
"The order by query does not have a corresponding composite index that it can be served from."
- 2) Using Indexing Policy (name -> descending, city -> ascending) works correctly, lists items.



So yes, people above have right.  
- ASC is default for name, so query is equivalent to:  
SELECT \* FROM People p ORDER BY p.name ASC, p.city DESC  
- "The composite index also supports an ORDER BY clause with the opposite order on all paths."  
<https://learn.microsoft.com/en-us/azure/cosmos-db/index-policy#order-by-queries-on-multiple-properties>  
upvoted 5 times



  **[Removed]** 1 year, 11 months ago  
No. It supports, but it is not required. You can have (ASC, ASC), (DESC, DESC), (ASC, DESC), (DESC, ASC). So the answer is Descending  
upvoted 2 times



  **neelkanths** 1 year, 2 months ago  
answer is not descending  
upvoted 2 times



  **GCMan**  4 years, 7 months ago  
"name" field should be marked ascending (default if not specified). It's mislabeled  
upvoted 35 times

  **cloud\_exam1** 4 years, 7 months ago  
I think so. The answer is correct, but the name field should be marked ascending(default).  
We can find an example in the following link.  
<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-indexing-policy?tabs=dotnetv2%2Cpythonv3>  
<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-manage-indexing-policy?tabs=dotnetv2%2Cpythonv3>  
upvoted 15 times

  **warchoon** 2 years, 4 months ago  
The client wants to walk with his back to the front. Who are we to tell him what to do?  
upvoted 1 times

  **scottmct** 4 years, 4 months ago  
NO. Box 2 is "ascending"  
See explanation here:  
<https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy#order-by-queries-on-multiple-properties>  
"The composite index also supports an ORDER BY clause with the \_\_opposite order on all paths\_\_."  
The table in the section also shows an example similar to this question.  
upvoted 31 times

  **Rockm0uld** 3 years, 3 months ago  
the table also seems to suggest DESC followed by ASC is not supported by composite index? Confused  
upvoted 2 times

  **Rockm0uld** 3 years, 3 months ago

Looking again i think it refers to " What it will not support is non-matching clauses. if ASC, ASC or DESC, DESC will not match our question." as mentioned by edengoforit.

Has to be opposites to make any sense.

upvoted 1 times

  **ArturKon** 2 years, 9 months ago

I agree with you. For example index on (A asc, B asc) works for queries with ORDER BY (A asc, B asc) and (A desc, B desc).

Not working with ORDER BY (A asc, B desc), (A desc, B asc) or even (B asc, A asc).

upvoted 1 times

  **[Removed]** 1 year, 11 months ago

No. Box 2 is Descending

upvoted 1 times

  **4bd3116** Most Recent 9 months, 4 weeks ago

```
{
  "indexingPolicy": {
    "compositeIndexes": [
      {
        "path": "/name",
        "order": "ascending"
      },
      {
        "path": "/city",
        "order": "descending"
      }
    ]
  }
}
```

upvoted 1 times

  **neelkanths** 1 year, 2 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

This questions came in 3 exams me and my friend gave....and don't get confused with answer

It is


Composite indexes

ascending

as some people in comment have explained..

descending doesn't go....scored very good marks based on these answers

upvoted 6 times

  **Jass1nonly** 1 year, 4 months ago

Composite Index Sample ORDER BY Query Supported by Composite Index?

(name ASC, age ASC) SELECT \* FROM c ORDER BY c.name ASC, c.age asc Yes


(name ASC, age ASC) SELECT \* FROM c ORDER BY c.age ASC, c.name asc No

(name ASC, age ASC) SELECT \* FROM c ORDER BY c.name DESC, c.age DESC Yes

The 1st and 3rd lines have opposite order-by value, but the composite index (name ASC, age ASC) still supports them.

The value for city should be ascending.

upvoted 1 times

  **Stann07** 1 year, 6 months ago

Got that question on my exam DEC 18. went with given answer. scored 842. CAsE Study : Van Arsdel inc.

upvoted 1 times

  **megp123** 1 year, 7 months ago

The only way to confirm the answer is to try it out.

"ascending" is the correct answer.

"descending" - will result to "Message: {"Errors":["The order by query does not have a corresponding composite index that it can be served "]}"  
upvoted 3 times

🗨️ 👤 **bluetopp** 1 year, 7 months ago

Just tested with proposed answer. If you run the query you will get the following error:

```
{"Errors":["The order by query does not have a corresponding composite index that it can be served from."]}
```

If you instead use "ascending" as the top voted answer here suggest, the query passes.

So correct answer is:

1. compositeIndexes

2. ascending

upvoted 4 times

🗨️ 👤 **Vukile** 1 year, 7 months ago

On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors

upvoted 2 times

🗨️ 👤 **AndySmith** 1 year, 7 months ago

On exam 3-Nov-2023. Selected:

1) compositeIndexes

2) descending

upvoted 2 times

🗨️ 👤 **p2006** 1 year, 9 months ago

Got on 9/25/2023

compositeindexes

ascending

upvoted 2 times

🗨️ 👤 **leviata** 1 year, 9 months ago

Got this question in exam - 2023.09.25. Got Case Study Contoso

upvoted 1 times

🗨️ 👤 **mihailos** 1 year, 10 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 1 times

🗨️ 👤 **basquiat** 1 year, 10 months ago

Got this one on 2023-08-08

upvoted 1 times

🗨️ 👤 **applepie** 1 year, 11 months ago

got this question today, answer compositeIndex, descending - 7/30/2023, score 895/1000

upvoted 3 times

🗨️ 👤 **NightshadeRC** 1 year, 11 months ago

Had this question today: 2023-07-26

upvoted 1 times

🗨️ 👤 **juanckar** 1 year, 11 months ago

This was on the exam (July 2023). Went with proposed (city should be "asc"). Scored 917

upvoted 3 times

**HOTSPOT -**

You are building a traffic monitoring system that monitors traffic along six highways. The system produces time series analysis-based reports for each highway.

Data from traffic sensors are stored in Azure Event Hub.

Traffic data is consumed by four departments. Each department has an Azure Web App that displays the time series-based reports and contains a WebJob that processes the incoming data from Event Hub. All Web Apps run on App Service Plans with three instances.

Data throughput must be maximized. Latency must be minimized.

You need to implement the Azure Event Hub.

Which settings should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Setting	Value
Number of partitions	<div><div></div><div>3</div><div>4</div><div>6</div><div>12</div></div>
Partition Key	<div><div></div><div>Highway</div><div>Department</div><div>Timestamp</div><div>VM name</div></div>

**Answer Area**

Setting	Value
Number of partitions	<div><div></div><div>3</div><div>4</div><div>6</div><div>12</div></div>
Partition Key	<div><div></div><div>Highway</div><div>Department</div><div>Timestamp</div><div>VM name</div></div>

**Suggested Answer:**

Box 1: 6 -

The number of partitions is specified at creation and must be between 2 and 32.

There are 6 highways.

Box 2: Highway -

Reference:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

 **dantest**  4 years, 6 months ago

Partitions relate to producers - and the logical way to partition the incoming data is by the only value you have at that point, the highway name/id. So the selected answer is correct (6 Partitions, by Highway).

People are getting confused by the departments which would actually each be an event consumer with an associated Consumer Group which would have it's own isolated view of each of the highway partitions.

upvoted 139 times

 **MiraA** 3 years, 8 months ago



One thought...

The assignment mentions "Partition Key" (not "Partition Id").

"Producers can provide a value for the event key. When they do, a hashing-based partitioner determines a hash value from the key. The event then goes to the partition associated with that hash value."

<https://docs.microsoft.com/en-us/azure/architecture/reference-architectures/event-hubs/partitioning-in-event-hubs-and-kafka#distribute-events-to-partitions>



So it is possible for 2 or more highways to be - by a chance - hashed to a single partition leaving 1 or more partitions idle at all. If really unlucky then all 6 highways would be hashed to one partition.

The very same situation is with Highway, Department and VM name used as the "Partition Key" as they are discrete values (6 highways, 4 departments, N virtual machines). The Timestamp could do a better job - but "spraying" data from one highway across all partitions...

The best job could do "Partition Id".



"Producers can specify a partition ID with an event. The event then goes to the partition with that ID." So each highway could get hardcoded its own partition preserving order of the data.

upvoted 2 times

  **PhilLI** 3 years, 7 months ago

I guess you're right, although losing maximum availability.

upvoted 1 times

  **Dani\_ac7** 2 years, 10 months ago

If we have 4 groups of consumers and each consumer group, according to the theory, reads independently and isolated from the rest, highways and 6 partitions makes a lot of sense



upvoted 2 times

  **Robert12345Robert** 3 years, 12 months ago

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-faq>

The number of partitions in an event hub directly relates to the number of concurrent readers you expect to have

upvoted 9 times

  **PhilLI** 3 years, 7 months ago

If you read the features page you'll understand that this is not the simple answer to the question. <https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

Event receivers organised in consumer groups read all partitions. Via a leasing mechanism the receiver within a consumer group can make sure they don't read duplicate data.

Partitions are more related to Event producers (can be 1-on-1)

upvoted 2 times

  **minaritochuck** 3 years, 4 months ago



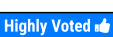
it's a best practice for publishers(producers) to remain unaware of the specific partitioning model chosen for an event hub and to only specify a partition key that is used to consistently assign related events to the same partition.

upvoted 2 times

  **minaritochuck** 3 years, 4 months ago

oh nevermind, I guess it would be most effective when partitions num matches the producers num.

upvoted 1 times

  **sasisang**  4 years, 7 months ago

The answer should be 4 and Highway.

Exam Topics - Please provide correct answers. What is the use of buying questions on your site...if you are not sure of the answer yourself

upvoted 43 times

  **kwaazaar** 4 years, 2 months ago

But theres 6 highways, so why not 6 partitions?

upvoted 2 times

  **clarionprogrammer** 4 years, 2 months ago

There are 4 consumers.

Partitions are a data organization mechanism that relates to the downstream parallelism required in consuming applications. The number of

partitions in an event hub directly relates to the number of concurrent readers you expect to have.

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-faq>

upvoted 2 times

  **guardna** 4 years, 1 month ago

But each department has 3 instances of the Web App/Job. So isn't there 12 consumers?

upvoted 7 times

  **warchoon** 2 years, 4 months ago

"The event publisher is only aware of its partition key, not the partition to which the events are published. This decoupling of key and partition insulates the sender from needing to know too much about the downstream processing. A per-device or user unique identity makes a good partition key, but other attributes such as geography can also be used to group related events into a single partition.



Specifying a partition key enables keeping related events together in the same partition and in the exact order in which they arrived. The partition key is some string that is derived from your application context and identifies the interrelationship of the events. A sequence of events identified by a partition key is a stream. A partition is a multiplexed log store for many such streams."

upvoted 1 times

  **jessicazheng** 3 years, 5 months ago



<https://www.linkedin.com/pulse/azure-event-hub-understanding-designing-partitions-unit-kamal-pathak>

upvoted 1 times

  **BrettusMaximus** 4 years, 2 months ago

There are 6 highways and 6 reports. Each department only needs to read one partition to produce their report on that one highway. If you had 4 partitions you would have to duplicate all the data 4 times (BAD) and then each department would need to read all the data and filter the data for one report (VERY BAD).

upvoted 11 times

  **BrettusMaximus** 4 years, 1 month ago

When you create the Hub it does not know either the Departments or the VMs

upvoted 3 times

  **Iaminall**  2 months ago

Number of Partitions: 12

Partition Key: highway\_id

You have 4 departments, each with a Web App running on 3 instances, which means 12 potential concurrent consumer instances ( $4 \times 3 = 12$ ).

upvoted 1 times

  **8c64e49** 3 months, 2 weeks ago

Number of Partitions (6):

Each partition in an Azure Event Hub operates as an independent channel for processing data. By matching the number of partitions to your six highways, each highway's data is routed to its own dedicated partition.

This parallelization ensures that data streams for different highways do not interfere with each other, maximizing throughput and reducing latency.

Partition Key ("Highway"):

The partition key determines how data is distributed across partitions. By using "Highway" as the key, all data related to a specific highway is consistently routed to the same partition.

This organization is critical for scenarios requiring ordered processing of data, like traffic pattern analysis or highway-specific monitoring.

If you were to use fewer partitions, multiple highways might share a partition, potentially creating bottlenecks. Conversely, more partitions than highways would be inefficient, as some partitions would remain underutilized.

upvoted 1 times

  **wafa\_chaari** 5 months ago

6 and highway,

departments are consumer groups.

(NB: you can use up to 5 consumers per partition where all receivers get all the events of the partition)

upvoted 1 times

🗨️ **cmmr** 7 months, 1 week ago

After spending a lot of time here, I'm going with a different answer:

- 3

- Highway

Although there are 4 departments, each department is part of a different consumer group, consuming messages in parallel. Each consumer group has 3 consumers, which will be assigned partitions.

According to MS: "We recommend that there's only one active receiver on a partition within a consumer group. However, in certain scenarios, you can use up to five consumers or receivers per partition where all receivers get all the events of the partition. If you have multiple readers on the same partition, then you process duplicate events. You need to handle it in your code, which isn't trivial. However, it's a valid approach in some scenarios."

Link: <https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

upvoted 1 times

🗨️ **Vichu\_1607** 7 months, 3 weeks ago

Recommended Settings:

Number of Partitions: 12

Partition Key: Highway identifier (e.g., highway\_id)

Number of Partitions:

The number of partitions should be chosen based on the expected throughput and the number of concurrent consumers. Since you have four departments, each with a WebJob processing the data, and each WebJob runs on three instances, you have a total of  $4 * 3 = 12$  consumers.

upvoted 2 times

🗨️ **Mattt** 8 months, 1 week ago

I gave the whole discussions to ChatGPT and here is the summery:

The final answer, based on the consensus and reasoning in the discussion, appears to be:

Number of partitions: 6

Partition key: Highway

The rationale is that since there are 6 highways, partitioning by highway ensures that events related to each highway are processed in order and kept separate. This structure supports scalability and ordered event processing, which is essential for time series analysis in the system.

upvoted 1 times

🗨️ **Christian\_garcia\_martin** 11 months, 1 week ago

copilot : 6 and Highway

upvoted 2 times

🗨️ **Hendrikdb** 1 year, 4 months ago

If your partition key is highway, no way your number can be 12. Only when your partition key is department, you can put 12 instances.

Even chatgpt and bard are proposing both aproaches but propose scaling it to the consumer side.

After long consideration, I sitl think highway and 6 instances is correct.

upvoted 1 times

🗨️ **11\_NickName\_11** 1 year, 7 months ago

When a client application sends events to an event hub without specifying a partition, events are automatically distributed among partitions in your event hub. If a partition isn't available for some reason, events are distributed among the remaining partitions. This behavior allows for the greatest amount of up time. For use cases that require the maximum up time, this model is preferred instead of sending events to a specific partition.

So maybe 12 & VM?

<https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-availability-and-consistency?tabs=dotnet#availability>

upvoted 1 times

🗨️ **AlbertoBT** 1 year, 7 months ago

Partition key:

Highway (Partitions relate to producers) You should not base your data on the architecture this is not scalable

Partitions: (Hub partitions not data partition this last are implicit by the Highway Partition key)

12

Partitioning allows for multiple parallel logs to be used for the same event hub and therefore multiplying the available raw IO throughput capacity.  
<https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features#advantages-of-using-partitions>

#### Mapping of events to partitions

The event publisher is only aware of its partition key, not the partition to which the events are published. This decoupling of key and partition insulates the sender from needing to know too much about the downstream processing. A per-device or user unique identity makes a good partition key...

<https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features#number-of-partitions>

Then in each VM the webjob can split the processes in batches using the Highway to speed up the query.

upvoted 1 times

🗨️ 👤 **adilkhan** 2 years, 3 months ago

6 and highway chatGPT

upvoted 1 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

Lol , i asked ChatGPT and it said 12. 4 departments \* 3 instances = 12 partitionkey needed

upvoted 1 times

🗨️ 👤 **RealRaymond** 2 years, 4 months ago

12 (4\*3) physical partitions; timestamp as partition key

upvoted 1 times

🗨️ 👤 **dimsok** 2 years, 4 months ago

The number of partitions to use in Azure Event Hub depends on the desired level of parallelism and the expected ingress rate. For maximum data throughput, you would want to use as many partitions as possible. Each partition allows for a separate stream of events to be processed in parallel, increasing the ingress rate. On the other hand, having too many partitions can lead to increased latency, as the events must be divided among the partitions

Answer is 12.

The partition key is a value that determines which partition an event should be sent to. The events with the same partition key are guaranteed to be ordered and processed by the same partition. In a traffic monitoring system, the partition key could be related to the highway being monitored.

Answer is highway.

upvoted 1 times

🗨️ 👤 **michaelfastov** 2 years, 7 months ago

<https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

"Your own applications must be able to keep up with processing the volume of events that are being sent into an event hub"

In our case events are highway reports (6). 4 applications are consumers.

6 is correct.

upvoted 1 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

selected answer is correct

upvoted 1 times

## DRAG DROP -

You are developing a microservices solution. You plan to deploy the solution to a multinode Azure Kubernetes Service (AKS) cluster.

You need to deploy a solution that includes the following features:

- ⇒ reverse proxy capabilities
- ⇒ configurable traffic routing
- ⇒ TLS termination with a custom certificate

Which components should you use? To answer, drag the appropriate components to the correct requirements. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

### Answer Area

Components	Action	Component
Helm		
Draft	Deploy solution.	
Brigade	View cluster and external IP addressing.	
KubeCtl	Implement a single, public IP endpoint that is routed to multiple microservices.	
Ingress Controller		
CoreDNS		
Virtual Kubelet		

### Answer Area

Suggested Answer:

Components	Action	Component
Helm		
Draft	Deploy solution.	Helm
Brigade	View cluster and external IP addressing.	KubeCtl
KubeCtl	Implement a single, public IP endpoint that is routed to multiple microservices.	Ingress Controller
Ingress Controller		
CoreDNS		
Virtual Kubelet		

Box 1: Helm -

To create the ingress controller, use Helm to install nginx-ingress.

Box 2: kubectl -

To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

### Box 3: Ingress Controller -

An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual Kubernetes services.

Incorrect Answers:

Virtual Kubelet: Virtual Kubelet is an open-source Kubernetes kubelet implementation that masquerades as a kubelet. This allows Kubernetes nodes to be backed by Virtual Kubelet providers such as serverless cloud container platforms.

CoreDNS: CoreDNS is a flexible, extensible DNS server that can serve as the Kubernetes cluster DNS. Like Kubernetes, the CoreDNS project is hosted by the

CNCF.

Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic> <https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking>

  **mlantonis** Highly Voted 3 years, 7 months ago

Box 1: Helm

Helm helps you manage Kubernetes applications — Helm Charts help you define, install, and upgrade even the most complex Kubernetes application. To create the ingress controller, use Helm to install nginx-ingress.

Box 2: Kubectl

The Kubernetes command-line tool, kubectl, allows you to run commands against Kubernetes clusters. To find the cluster IP address of a Kubernetes pod, use the kubectl get pod command on your local machine, with the option -o wide .

Box 3: Ingress Controller

An ingress controller is a piece of software that provides reverse proxy, configurable traffic routing, and TLS termination for Kubernetes services. Kubernetes ingress resources are used to configure the ingress rules and routes for individual Kubernetes services. Using an ingress controller and ingress rules, a single IP address can be used to route traffic to multiple services in a Kubernetes cluster.

upvoted 76 times

  **mlantonis** 3 years, 7 months ago

Reference:

<https://helm.sh>

<https://kubernetes.io/docs/tasks/tools>

<https://kubernetes.io/docs/concepts/services-networking/ingress-controllers>

<https://docs.microsoft.com/bs-cyrl-ba/azure/aks/ingress-basic>



<https://www.digitalocean.com/community/tutorials/how-to-inspect-kubernetes-networking>

upvoted 9 times

  **robin1337** 2 years, 8 months ago

Correct, I am a certified CKA und CKAD. The question isn't that well written though as you could also deploy a solution with kubectl...


upvoted 11 times

  **Leandromellor** Highly Voted 4 years, 1 month ago

I believe there is no AKS question in the exam:

<https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B>

upvoted 19 times

  **matejka** 4 years, 1 month ago

Haven't done the exam yet, but found many comments (eg. on Udemy tutorials) that Kuberentes actually IS part of the exam.

upvoted 6 times

  **luppittegui** 4 years ago

It's not: <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4oZ7B>

"Azure Kubernetes Service (AKS) is out of scope"

upvoted 4 times

  **Cornholioz** 4 years ago

And yet, here we are with this question. I'm guessing it has appeared in the exam and hence here. I am wondering what's actually "excluded" from Azure in this weird exam. Can't say it for many, but this one is a poorly constructed exam.

upvoted 7 times

  **daveonplanetearth** 3 years, 10 months ago



I also thought AKS was excluded but people say they have had AKS questions in the exam. I don't understand why they would want to exclude it.

upvoted 1 times

  **Rockm0uld** 2 years, 9 months ago

I'm pretty sure I got this question on the exam, they be messin' with us

upvoted 1 times

  **cherylm** 3 years, 10 months ago

I've also seen a lot of people mentioned in facebook groups they got kubernetes questions on the real exam az-204

upvoted 4 times

  **DaveOla**  9 months, 3 weeks ago



Kubernetes is no longer covered in the exam

upvoted 3 times

  **bgbgvfvf** 1 year ago

It is correct.

upvoted 1 times

  **wrak** 1 year, 10 months ago

out-of-the-scope for the AZ-204 Exam; please remove from the Question Bank

upvoted 7 times

  **macobuzi** 1 year, 4 months ago

Microsoft doesn't give a f about the scope. Better safe than sorry!

upvoted 4 times

  **nvtienanh** 2 years, 3 months ago

Box 1: Kubectl or Helm

upvoted 2 times

  **azurelearner666** 3 years, 6 months ago

It is correct.

Helm

KubeCtl

Ingress


simple as that.

upvoted 5 times

  **mlantonis** 3 years, 7 months ago

Answer seems legit, but it's weird getting Kubernetes questions, when AKS is out of the scope of the exam.

upvoted 8 times

  **glam** 3 years, 7 months ago

correct.

upvoted 1 times

  **Frakandel** 3 years, 7 months ago


Answer:

1. KubeCtl or Helm

2. KubeCtl

3. Ingress Controller

upvoted 4 times

  **Annamarie0408** 3 years, 8 months ago

Just took this test last week (last week of March) AKS -- WAS IN FACT -- In the test.

upvoted 7 times

  **i\_ra** 3 years, 8 months ago



Is these questions still valid? I'm taking the exam this week can I depend on these?? please answer me

upvoted 2 times

  **kwaazaar** 3 years, 9 months ago

helm and kubectl can both be used to deploy.

upvoted 4 times

  **svaza** 3 years, 10 months ago

Answer is right

upvoted 2 times



DRAG DROP -

You are implementing an order processing system. A point of sale application publishes orders to topics in an Azure Service Bus queue. The Label property for the topic includes the following data:

Property	Description
ShipLocation	the country/region where the order will be shipped
CorrelationId	a priority value for the order
Quantity	a user-defined field that stores the quantity of items in an order
AuditedAt	a user-defined field that records the date an order is audited

The system has the following requirements for subscriptions:

Subscription type	Comments
FutureOrders	This subscription is reserved for future use and must not receive any orders
HighPriorityOrders	Handle all high priority orders and international orders
InternationalOrders	Handle orders where the country/region is not United States
HighQuantityOrders	Handle only orders with quantities greater than 100 units
AllOrders	This subscription is used for auditing purposes. This subscription must receive every single order. AllOrders has an Action defined that updates the AuditedAt property to include the date and time it was received by the subscription.

You need to implement filtering and maximize throughput while evaluating filters.

Which filter types should you implement? To answer, drag the appropriate filter types to the correct subscriptions. Each filter type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

### Filter types

SQLFilter

CorrelationFilter

No Filter

### Answer Area

#### Subscription

FutureOrders

HighPriorityOrders

InternationalOrders

HighQuantityOrders

AllOrders

#### Filter type

filter type

filter type

filter type

filter type

filter type

Suggested Answer:

### Filter types

SQLFilter

CorrelationFilter

No Filter

### Answer Area

#### Subscription

FutureOrders

HighPriorityOrders

InternationalOrders

HighQuantityOrders

AllOrders

#### Filter type

SQLFilter

CorrelationFilter

SQLFilter

SQLFilter

No Filter

FutureOrders: SQLFilter -

HighPriorityOrders: CorrelationFilter

CorrelationID only -

InternationalOrders: SQLFilter -

Country NOT USA requires an SQL Filter

HighQuantityOrders: SQLFilter -

Need to use relational operators so an SQL Filter is needed.

AllOrders: No Filter -

SQL Filter: SQL Filters - A SqlFilter holds a SQL-like conditional expression that is evaluated in the broker against the arriving messages' user-defined properties and system properties. All system properties must be prefixed with sys. in the conditional expression. The SQL-language subset for filter conditions tests for the existence of properties (EXISTS), as well as for null-values (IS NULL), logical NOT/AND/OR, relational operators, simple numeric arithmetic, and simple text pattern matching with LIKE.

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against ContentType, Label, MessageId, ReplyTo,

ReplyToSessionId, SessionId, To, and any user-defined properties. A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical

AND condition, meaning for the filter to match, all conditions must match.

Boolean filters - The TrueFilter and FalseFilter either cause all arriving messages (true) or none of the arriving messages (false) to be selected for the subscription.

Reference:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

  **stylebc**  4 years, 7 months ago

I think that it should be

-Correlation Filter (with the not existing value of any field to avoid getting any message)

-SQL filter (as we need to get all high priority AND international orders, but for Correlation filter: A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter and we need not equal)

-SQL filter

-SQL filter

-No Filter

upvoted 147 times

  **ferut** 4 years, 1 month ago

FutureOrder, not based on the property -> Correlation Filter.

AllOrders, it's clear.. no filter.

The rest is based on one or more properties -> SQL filter

I agree with @stylebc

upvoted 16 times

  **Mal22002** 3 years, 1 month ago

How do you justify No Filter for All Orders when you need to take action?

Each newly created topic subscription has an initial default subscription rule. If you don't explicitly specify a filter condition for the rule, the applied filter is the true filter that enables all messages to be selected into the subscription. The default rule has no associated annotation action

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

upvoted 1 times

  **Mal22002** 3 years, 1 month ago

Check this example with Action which is a requirement for AllOrders.

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-filter-examples#net-example-for-creating-subscription-filters>

upvoted 1 times


  **Mal22002** 3 years, 1 month ago

FutureOrders should be SQLFilter as you can check with the condition on AuditedAt user's property or EnqueuedTimeUtc system property but that is required a greater than condition but correlation filter condition needs to be matched.

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties.

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

upvoted 3 times

  **Mal22002** 3 years, 1 month ago

HighPriorityOrders: CorrelationFilter

You can apply to multiple system or user-defined properties and when multiple properties then filter combined them with AND logical operator.

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

Correlation Filters - A CorrelationFilter holds a set of conditions that are matched against one or more of an arriving message's user and system properties. A common use is to match against the CorrelationId property, but the application can also choose to match against the following properties:

ContentType

Label

MessageId

ReplyTo

ReplyToSessionId

SessionId

To

any user-defined properties.

A match exists when an arriving message's value for a property is equal to the value specified in the correlation filter. For string expressions, the comparison is case-sensitive. When specifying multiple match properties, the filter combines them as a logical AND condition, meaning for the filter to match, all conditions must match.

upvoted 2 times

  **TEMPKAKAM**  4 years, 7 months ago

The Correct answers are:

No Filter

Correlation Filter

SQL filter

SQL filter



SQL filter

upvoted 40 times

  **Christian\_garcia\_martin** 11 months, 1 week ago

no sense at all , you want future filters and do nothing? , you want all the orders and you apply filter?

upvoted 1 times

  **rashjan** 4 years, 7 months ago

Why a filter for all orders?

upvoted 3 times

  **MrNair007** 4 years, 7 months ago

See the basic is subscription does not have filter then does not receive any data. And since all orders needs all the orders it should have SQL filter with 1=1 so that all orders are passed to it.

upvoted 2 times

  **brtx** 4 years, 7 months ago

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

"If you don't explicitly specify a filter condition for the rule, the applied filter is the true filter that enables all messages to be selected into the subscription."

upvoted 25 times

  **danielcr** 4 years, 4 months ago

So.. The answer is correct.

upvoted 3 times

  **TakumaK** 4 years, 1 month ago

if no filter explicitly specified, the true filter will be assigned which enables ALL messages. So your comment is wrong and makes confused.

upvoted 11 times

  **JustHereToLearn** 2 years, 6 months ago

I somewhat agree with this answer after going through lot of resources. The only thing that bothers me is answer for the first and last question.



When you create a subscription by default a SQL filter is added which is 1=1 i.e. always evaluates to true.

So my source of confusion is how "No Filter" is treated. I have two interpretations either "No Filter" is we don't add any filter or the other is

absence of a filter.

Mostly I believe it is treated as absence of a filter in which case the this answer looks correct.

upvoted 2 times

  **cbn** 4 years, 4 months ago

How do you justify Correlation Filter for HighPriorityOrders, when it involves Region <> US ?

I think it should be SQL filter as well.

upvoted 14 times

  **vizay**  2 months ago

So, what's the best option to block everything?

Use an SQLFilter like:

$1 = 0$



This is always false, guaranteed – and it will never allow any messages through.

upvoted 1 times

  **wafa\_chaari** 5 months ago

Given answer is correct

upvoted 1 times

  **Mattt** 8 months, 1 week ago

The given answer is correct

upvoted 1 times

  **ns4098** 9 months, 2 weeks ago

Final Answer:

FutureOrders: SQLFilter

HighPriorityOrders: CorrelationFilter

InternationalOrders: SQLFilter

HighQuantityOrders: SQLFilter

AllOrders: No Filter

upvoted 2 times

  **Christian\_garcia\_martin** 11 months, 1 week ago

copilot :

for futures ---> correlation,

highpriority , internal , hightQuantity -----> SQLFilter

all -----> no filter

upvoted 1 times

  **gebpt** 1 year ago

1) SQL  $1=0$

2) SQL

3) SQL

4) SQL

5) No filter - takes all

upvoted 2 times

  **neelkanths** 1 year, 2 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

went with below answer based on the example given on microsoft learn....go with it...scored very good marks may be highest with these answers

-SQL filter

-Correlation Filter

-SQL filter

-SQL filter

-No Filter

upvoted 7 times

  **jobolesonihal** 1 year, 3 months ago

Given answer is correct.

upvoted 2 times

🗨️ 👤 **raymond\_abcd** 1 year, 5 months ago

Missing the Boolean filter option in the possible answers. Because that would be the one to choose for the future order rule. Otherwise I think you can choose SQL filter or correlation filter for the future order by making the condition always filter out all messages.

The other ones I agree with the given answers.

upvoted 1 times

🗨️ 👤 **emysa341** 1 year, 6 months ago

- No Filter

- Correlation

- Correlation

- SQL Filter

- SQL Filter

upvoted 1 times

🗨️ 👤 **lednari** 1 year, 7 months ago

Just tested this in azure portal

- No filter (if default sql 1=1 filter is removed) results in no messages

- sql filter 1=0 can be used for same purpose

- sql filter 1=1 will match everything and is the default when adding a subscription

Therefore:

No filter

Sql (need to match hi pri AND international orders)

Sql

Sql

Sql (default 1=1)

upvoted 4 times

🗨️ 👤 **leviatas** 1 year, 9 months ago

Got this quesiton in examn, went with answer. - 2023.09.25. Got Case Study Contoso

upvoted 3 times

🗨️ 👤 **longnguyendh** 1 year, 11 months ago

It should be SQL, Correlation, SQL, SQL and NO Filter

upvoted 2 times

🗨️ 👤 **applepie** 1 year, 11 months ago

got this question today, go with answer -

Correlation

Correlation

SQL

SQL

No filter

7/30/2023, score 895/1000

upvoted 2 times

🗨️ 👤 **NightshadeRC** 1 year, 11 months ago

Had this question today: 2023-07-26

I went:

- Correlation Filter

- SQL Filter

- SQL Filter

- SQL Filter

- No Filter

Every question on the exam was on ExamTopics. I entered every recommended answer and got 940

upvoted 10 times

🗨️ 👤 **Ciupaz** 1 year, 7 months ago

How do you remember all 375 questions?

upvoted 4 times

## DRAG DROP -

Your company has several websites that use a company logo image. You use Azure Content Delivery Network (CDN) to store the static image. You need to determine the correct process of how the CDN and the Point of Presence (POP) server will distribute the image and list the items in the correct order.

In which order do the actions occur? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order. Select and Place:

**Actions****Answer Area**

If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.

A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.

Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the file from cache if the TTL has not expired.

The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.

**Suggested Answer:****Actions**

If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.

A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.

Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the file from cache if the TTL has not expired.

The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.

**Answer Area**

A user requests the image from the CDN URL. The DNS routes the request to the best performing POP location.

If no edge servers in the POP have the image in cache, the POP requests the file from the origin server.

The origin server returns the logo image to an edge server in the POP. An edge server in the POP caches the logo image and returns the image to the client.

Subsequent requests for the file may be directed to the same POP using the CDN logo image URL. The POP edge server returns the file from cache if the TTL has not expired.



Step 1: A user requests the image..

A user requests a file (also called an asset) by using a URL with a special domain name, such as <endpoint name>.azureedge.net. This name can be an endpoint hostname or a custom domain. The DNS routes the request to the best performing POP location, which is usually the POP that is geographically closest to the user.

Step 2: If no edge servers in the POP have the..

If no edge servers in the POP have the file in their cache, the POP requests the file from the origin server. The origin server can be an Azure Web App, Azure

Cloud Service, Azure Storage account, or any publicly accessible web server.

Step 3: The origin server returns the..

The origin server returns the file to an edge server in the POP.

An edge server in the POP caches the file and returns the file to the original requestor (Alice). The file remains cached on the edge server in the POP until the time-to-live (TTL) specified by its HTTP headers expires. If the origin server didn't specify a TTL, the default TTL is seven days.

Step 4: Subsequent requests for..

Additional users can then request the same file by using the same URL that the original user used, and can also be directed to the same POP.

If the TTL for the file hasn't expired, the POP edge server returns the file directly from the cache. This process results in a faster, more responsive user experience.

Reference:

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

🗨️ 👤 **Ash111** Highly Voted 3 years, 7 months ago

Given ans is correct  
upvoted 96 times

🗨️ 👤 **TonyMel** 1 year, 3 months ago

correct, in 2023Mar24, score: 904/1000  
upvoted 6 times

🗨️ 👤 **KhabibcandefeatGSP** 3 years ago

It feels so good when there are no debates on what is the correct answer.  
upvoted 67 times

🗨️ 👤 **Tealon** Highly Voted 3 years, 6 months ago

The given answer is correct.  
upvoted 10 times

🗨️ 👤 **Tarajee** Most Recent 9 months, 3 weeks ago

Got on my exam 2023sept  
upvoted 2 times

🗨️ 👤 **warchoon** 1 year, 4 months ago

"... returns the file directly from the cache. This process results in a faster, more responsive user experience" - last in order.  
Hmmm. Fine  
upvoted 1 times

🗨️ 👤 **kyawtlearning** 1 year, 10 months ago

Given answer is correct.  
upvoted 2 times

🗨️ 👤 **Eltooth** 2 years ago

Answer is correct.  
upvoted 2 times

🗨️ 👤 **gursimran\_s** 2 years, 2 months ago

Got it in April 2023  
upvoted 2 times

🗨️ 👤 **OPT\_001122** 1 year, 7 months ago

time travel!!!  
upvoted 10 times

🗨️ 👤 **meoukg** 2 years, 3 months ago

Got it on 03/2022, chose the same as the given answer.  
upvoted 2 times

🗨️ 👤 **silva\_831** 3 years ago

The given answer is correct  
upvoted 2 times

🗨️ 👤 **mlantonis** 3 years, 1 month ago

The Answer is correct

Reference:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-overview>  
upvoted 5 times

🗨️ 👤 **warchoon** 1 year, 4 months ago

"How it works" does not exactly represent the order.  
"returns the file from cache if" should be before "if no ... image in cache"  
upvoted 1 times

🗨️ 👤 **warchoon** 1 year, 4 months ago

I know that "Subsequent" is the bad word for this choice and the author could be stupid enough to use wrong order.  
upvoted 1 times

🗨️ 👤 **prabhjot** 3 years, 1 month ago

correct ans

upvoted 1 times

🗨️ 👤 **UnknowMan** 3 years, 1 month ago

The given answer is correct.

upvoted 2 times

🗨️ 👤 **glam** 3 years, 1 month ago

correct.

upvoted 1 times

🗨️ 👤 **Fr3ddy** 3 years, 5 months ago

The az-204 exam voucher costs 100 euros.

This value includes Study material + dump

upvoted 5 times

🗨️ 👤 **azurelearner666** 3 years ago

And? this comment is silly and does not help.

It's like saying: the sky is blue, except when it is clouded (or dark).

Admin: can you delete this useless message above? (and mine)

upvoted 11 times

🗨️ 👤 **coffecold** 1 year, 8 months ago

Admin, can you delete all "this answer is correct" minus one, and not delete the comment of Fr3ddy?

upvoted 2 times

🗨️ 👤 **mmdex** 1 year, 4 months ago

Admin, can you delete this whole discussion? I don't see any comment that would have any value whatsoever.

upvoted 2 times

🗨️ 👤 **Saurabh\_Kulkarni** 2 years, 9 months ago

Lol...

upvoted 4 times

🗨️ 👤 **pac1311** 3 years, 5 months ago

correct!

upvoted 2 times

🗨️ 👤 **sndp** 3 years, 6 months ago

Answer is correct.

upvoted 5 times



You are developing an Azure Cosmos DB solution by using the Azure Cosmos DB SQL API. The data includes millions of documents. Each document may contain hundreds of properties.

The properties of the documents do not contain distinct values for partitioning. Azure Cosmos DB must scale individual containers in the database to meet the performance needs of the application by spreading the workload evenly across all partitions over time.

You need to select a partition key.

Which two partition keys can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. a single property value that does not appear frequently in the documents
- B. a value containing the collection name
- C. a single property value that appears frequently in the documents
- D. a concatenation of multiple property values with a random suffix appended
- E. a hash suffix appended to a property value

**Suggested Answer: DE**

You can form a partition key by concatenating multiple property values into a single artificial partitionKey property. These keys are referred to as synthetic keys.


Another possible strategy to distribute the workload more evenly is to append a random number at the end of the partition key value. When you distribute items in this way, you can perform parallel write operations across partitions.

Note: It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys>

Community vote distribution



**TEMPKAKAM** Highly Voted 4 years, 7 months ago

The given answer is correct

upvoted 74 times

**mlantonis** Highly Voted 4 years, 1 month ago

D and E

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/synthetic-partition-keys>

upvoted 18 times

**vizay** Most Recent 2 months ago

**Selected Answer: DE**

Use a custom key – like:

A combination of fields (e.g., userId + timestamp)

Or a random ID/hash to spread documents evenly

"A hash suffix appended to a property value" means:

Take a property value (like a userId), then add a random or hashed number to the end to make it more unique.

Partition key = "user123#1"

Partition key = "user123#2"

Partition key = "user123#3"

upvoted 1 times

🗨️ 👤 **wafa\_chaari** 5 months ago

**Selected Answer: DE**

you must construct a synthetic partition key. 1-by concatenating multiple property values into a single partitionkey (expl: name=azure , date= 2024--> partitionkey= azure2024)

2-use a partition key with a random suffix ( expl azure1 , azure3..)

upvoted 1 times

🗨️ 👤 **Vichu\_1607** 7 months, 3 weeks ago

**Selected Answer: DE**

D. a concatenation of multiple property values with a random suffix appended

E. a hash suffix appended to a property value

upvoted 1 times

🗨️ 👤 **jobolesonihal** 1 year, 3 months ago

**Selected Answer: CE**

Copilot says C and E

With D - random suffice for each item may be an overkill.

upvoted 1 times

🗨️ 👤 **JoseManel** 1 year, 7 months ago

**Selected Answer: DE**

D. a concatenation of multiple property values with a random suffix appended Most Voted

E. a hash suffix appended to a property value Most Voted

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 11 months ago

Given answer is incorrect. Correct answer: AE

upvoted 1 times

🗨️ 👤 **Net\_IT** 1 year, 8 months ago

That is not the given answer? It is D and E.

upvoted 1 times

🗨️ 👤 **narenazure** 2 years, 5 months ago

**Selected Answer: DE**

It's the best practice to have a partition key with many distinct values, such as hundreds or thousands. The goal is to distribute your data and workload evenly across the items associated with these partition key values. If such a property doesn't exist in your data, you can construct a synthetic partition key

upvoted 2 times

🗨️ 👤 **mdg3501** 2 years, 6 months ago

got this on 2022-12-7

upvoted 5 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: DE**

D. a concatenation of multiple property values with a random suffix appended Most Voted

E. a hash suffix appended to a property value Most Voted

upvoted 2 times

🗨️ 👤 **TheExamMaster2020** 2 years, 7 months ago

Did my exam on 15th November 2022. This question was on it.

upvoted 4 times

🗨️ 👤 **coffecold** 2 years, 8 months ago

at option E : a hash value will always deliver the same result on the same data. It is not a random value as stated in the proposed solution.

Since A,B and C fall off, and two answers must be chosen still I would go for D,E.

upvoted 2 times

🗨️ 👤 **xRiot007** 3 years ago

**Selected Answer: DE**

If no property in the document data will have unique values, you need to make one.

This is called a synthetic partition key.

These sorts of keys are made by adding a unique suffix at the end of some property.

One other way is to create a property that will have the hashed data + a random suffix.

The objective is to have a property that is random enough so that you can rely on it to be your key.

upvoted 6 times

🗨️ 👤 **Eltooth** 3 years ago

**Selected Answer: DE**

D & E are the correct answers.

upvoted 1 times

🗨️ 👤 **Rini100** 3 years, 2 months ago

Got this on 20 Apr 2022

upvoted 2 times

🗨️ 👤 **oescm** 3 years, 4 months ago

Got this one 02/2022. Went with highly voted answer.

upvoted 6 times

## HOTSPOT -

You are developing an Azure-hosted e-commerce web application. The application will use Azure Cosmos DB to store sales orders. You are using the latest SDK to manage the sales orders in the database.

You create a new Azure Cosmos DB instance. You include a valid endpoint and valid authorization key to an appSettings.json file in the code project.

You are evaluating the following application code: (Line number are included for reference only.)

```

01 using System;
02 using System.Threading.Tasks;
03 using Microsoft.Azure.Cosmos;
04 using Microsoft.Extensions.Configuration;
05 using Newtonsoft.Json;
06 namespace SalesOrders
07 {
08     public class SalesOrder
09     {
10         . . .
11     }
12     internal class ManageSalesOrders
13     {
14         private static async Task GenerateSalesOrders()
15         {
16             IConfigurationRoot configuration = new ConfigurationBuilder().AddJsonFile("appSettings.json").Build();
17             string endpoint = configuration["EndPointUrl"];
18             string authKey = configuration["AuthorizationKey"];
19             using CosmosClient client = new CosmosClient(endpoint, authKey);
20             Database database = null;
21             using (await client.GetDatabase("SalesOrders").DeleteStreamAsync()) { }
22             database = await client.CreateDatabaseIfNotExistsAsync("SalesOrders");
23             Container container1 = await database.CreateContainerAsync(id: "Container1", partitionKeyPath: "/AccountNumber");
24             Container container2 = await database.CreateContainerAsync(id: "Container2", partitionKeyPath: "/AccountNumber");
25             SalesOrder salesOrder1 = new SalesOrder() { AccountNumber = "123456" };
26             await container1.CreateItemAsync(salesOrder1, new PartitionKey(salesOrder1.AccountNumber));
27             SalesOrder salesOrder2 = new SalesOrder() { AccountNumber = "654321" };
28             await container1.CreateItemAsync(salesOrder2, new PartitionKey(salesOrder2.AccountNumber));
29             SalesOrder salesOrder3 = new SalesOrder() { AccountNumber = "109876" };
30             await container2.CreateItemAsync(salesOrder3, new PartitionKey(salesOrder3.AccountNumber));
31             _ = await database.CreateUserAsync("User1");
32             User user1 = database.GetUser("User1");
33             _ = await user1.ReadAsync();
34         }
35     }
36 }

```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
A database named SalesOrders is created. The database will include two containers.	<input type="radio"/>	<input type="radio"/>
Container1 will contain two items.	<input type="radio"/>	<input type="radio"/>
Container2 will contain one item.	<input type="radio"/>	<input type="radio"/>

## Answer Area

	Statements	Yes	No
Suggested Answer:	A database named SalesOrders is created. The database will include two containers.	<input checked="" type="radio"/>	<input type="radio"/>
	Container1 will contain two items.	<input checked="" type="radio"/>	<input type="radio"/>
	Container2 will contain one item.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes -

The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it.

The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service.

Box 2: Yes -

The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service.

Box 3: Yes -

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync> <https://docs.microsoft.com/en-us/dotnet/api/azure.cosmos.cosmoscontainer.createitemasync>

  **mlantonis** Highly Voted 4 years ago

Box 1: Yes

The createDatabaseIfNotExistsAsync method checks if a database exists, and if it doesn't, create it. (Line 22)

The Database.CreateContainerAsync method creates a container as an asynchronous operation in the Azure Cosmos service. (Line 23 and 24)



Box 2: Yes

The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service. (Line 26 and 28)

Box 3: Yes

The CosmosContainer.CreateItemAsync method creates an item as an asynchronous operation in the Azure Cosmos service. (Line 30)

upvoted 91 times

  **Esward** 2 years, 5 months ago

Agreed

upvoted 1 times

  **chettir01** 2 years, 6 months ago

The first box is not correct, the name of database is database and not salesOrder !

upvoted 1 times

  **chettir01** 2 years, 6 months ago

SalesOrders is the databaseID, not the name !

upvoted 1 times

  **chettir01** 2 years, 6 months ago

// New instance of Database class referencing the server-side database

// The name of instance is database2, and we need an ID to create it !

Database database2 = await client.CreateDatabaseIfNotExistsAsync(

id: "adventureworks-2"

);

upvoted 1 times

  **Vladimir\_Gajinov** 1 year, 9 months ago

You are right but in Azure Cosmos DB, the database name is typically the same as the database ID, but spaces in the ID are replaced with hyphens ("-").

upvoted 1 times

  **Vladimir\_Gajinov** 1 year, 9 months ago

So, this answer is correct one.

upvoted 1 times

  **mlantonis** 4 years ago

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient.createdatabaseifnotexistsasync>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.database.createcontainerasync>

<https://docs.microsoft.com/en-us/dotnet/api/azure.cosmos.cosmoscontainer.createitemasync>

upvoted 7 times

  **AOE** Highly Voted 3 years, 10 months ago



Line 21 is tricky, it assumes the database is already created as it calls DeleteStreamAsync. I'm confused.

upvoted 7 times

  **warchoon** 2 years, 4 months ago

It returns ResponseMessage which is IDisposable

upvoted 1 times

  **Knightie** 2 years, 10 months ago

// Delete a Database resource where database\_id is the ID property of the Database resource you wish to delete.

Database database = this.cosmosClient.GetDatabase(database\_id);

await database.DeleteStreamAsync();

That is deleted so the create if not exist will surely create the DB

upvoted 1 times

  **coffecold** 2 years, 8 months ago

"using" is equal to a "try-finally" with noting in finally so the null exception is already been taken care of .

<https://stackoverflow.com/questions/2522822/will-dispose-be-called-in-a-using-statement-with-a-null-object>

upvoted 2 times

  **ReyPirata** Most Recent 1 year, 10 months ago

On my exam 2023-08-20. Scored 925

Yes

Yes

Yes

upvoted 4 times

  **ks321** 1 year, 11 months ago

Box 1: Yes

Box 2: Yes

Box 3: No - container2.Create is called only once

upvoted 1 times

  **macobuzi** 1 year, 10 months ago



I think you are confused! It said container2 has 1 item. So box3 - Yes is correct

upvoted 3 times

  **Christian\_garcia\_martin** 11 months, 1 week ago

exactly

upvoted 2 times

  **Videira** 2 years, 4 months ago

On my exam 2023-02-25

upvoted 1 times

  **Lucifer14** 2 years, 6 months ago

In 28-12-2022 exam

upvoted 3 times

🗨️ 👤 **TheExamMaster2020** 2 years, 7 months ago

Did my exam on 15th November 2022. This question was on it.

upvoted 2 times

🗨️ 👤 **sam5678** 2 years, 8 months ago

yes yes yes

upvoted 1 times

🗨️ 👤 **chettir01** 2 years, 6 months ago

No, it's not correct, The first box is not correct, the name of database on the code is database and not salesOrder !

upvoted 1 times

🗨️ 👤 **chettir01** 2 years, 6 months ago

SalesOrders is the databaseID, not the name !

upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

Answer is correct:

Yes

Yes

Yes

upvoted 1 times

🗨️ 👤 **chettir01** 2 years, 6 months ago

No, it's not correct, The first box is not correct, the name of database on the code is database and not salesOrders !

SalesOrders is the ID of the database not it's name !

upvoted 1 times

🗨️ 👤 **AymanAkk** 2 years, 6 months ago

there is no name of database, that s a variable called database of type Database and it will receive the newly created database with name salesOrder !

upvoted 4 times

🗨️ 👤 **nargzul** 4 years ago

I'm just curious, how do you know that Container 1 contains 2 items and container2 contains 1 item and not the opposite? I see that we are partitioning on the account number, but not sure to understand how the partition is made?

upvoted 2 times

🗨️ 👤 **BroGood** 4 years ago

see Lines 26 & 28 contain items for Container 1

Only Line 30 contains an item for Container 2

upvoted 5 times

🗨️ 👤 **ariel\_dev** 4 years ago

container1 is called two times adding the items, container2 just one.

upvoted 6 times

🗨️ 👤 **Knightie** 2 years, 10 months ago

container 1 will have 2 partitions, container 2 only 1 partition because partition key is the account number. But container 1 will still have 2 items, 1 item per partition. Hope this help.

upvoted 1 times

🗨️ 👤 **AlokSingh** 4 years ago

Correct Answer

upvoted 1 times

## DRAG DROP -

You develop an Azure solution that uses Cosmos DB.

The current Cosmos DB container must be replicated and must use a partition key that is optimized for queries.

You need to implement a change feed processor solution.

Which change feed processor components should you use? To answer, drag the appropriate components to the correct requirements. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view the content.

NOTE: Each correct selection is worth one point.

Select and Place:

**Components**

Host

Delegate

Lease container

Monitored container

**Answer Area****Requirement**

Store the data from which the change feed is generated.

Coordinate processing of the change feed across multiple workers.

Use the change feed processor to listen for changes.

Handle each batch of changes.

**Component**

Component

Component

Component

Component

**Suggested Answer:****Components**
  
  
  

**Answer Area****Requirement**

Store the data from which the change feed is generated.

Coordinate processing of the change feed across multiple workers.

Use the change feed processor to listen for changes.

Handle each batch of changes.

**Component**

Monitored container

Lease container

Host

Delegate

Box 1: The monitored container -

The monitored container has the data from which the change feed is generated. Any inserts and updates to the monitored container are reflected in the change feed of the container.

Box 2: The lease container -

The lease container acts as a state storage and coordinates processing the change feed across multiple workers. The lease container can be stored in the same account as the monitored container or in a separate account.

Box 3: The host: A host is an application instance that uses the change feed processor to listen for changes. Multiple instances with the same lease configuration can run in parallel, but each instance should have a different instance name.

Box 4: The delegate -

The delegate is the code that defines what you, the developer, want to do with each batch of changes that the change feed processor reads.


Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-processor>

 **MattXu** Highly Voted 4 years ago


The given answer is correct.

upvoted 60 times

 **TakumaK** 3 years, 12 months ago

You are right. The given answer is correct.

upvoted 11 times

 **john4p** 3 years, 6 months ago

And your given answer, that MattXu is right, is also correct.





upvoted 19 times

  **dbobspurfpo** 3 years, 6 months ago


that observation is correct, john

upvoted 15 times

  **Stann07** 1 year, 7 months ago



dbobspurfpo, you are right to say that the observation of john is correct.

upvoted 2 times

  **gmishra88**  2 years, 8 months ago

Microsoft has obscure names in documentation and has nothing to do with the azure component itself. "Delegate", "Host component", "Compute instance". How's remembering this from one page they wrote mean anything. We do not call any of them these when we implement the change feed processor

upvoted 15 times

  **vizay**  2 months ago

☞ Think of it like this:

Imagine you're watching a door for packages ☞:

Source container = The place where packages (data) arrive

Lease container = A log of which packages you already saw

Delegate = Your job (like moving the package somewhere or processing it)

Host = The person doing the work

upvoted 2 times

  **wafa\_chaari** 5 months ago

correct

The change feed processor has four main components:

The monitored container: The monitored container has the data from which the change feed is generated. Any inserts and updates to the monitored container are reflected in the change feed of the container.

The lease container: The lease container acts as state storage and coordinates the processing of the change feed across multiple workers. The lease container can be stored in the same account as the monitored container or in a separate account.

The compute instance(host): A compute instance hosts the change feed processor to listen for changes. Depending on the platform, it might be represented by a virtual machine (VM), a Kubernetes pod, an Azure App Service instance, or an actual physical machine. The compute instance has a unique identifier that's called the instance name throughout this article.

The delegate: The delegate is the code that defines what you, the developer, want to do with each batch of changes that the change feed processor reads.

upvoted 1 times

  **royalbaby** 11 months, 2 weeks ago



Got this Q on 07/07. Went with given answer.

upvoted 1 times

  **Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 3 times

  **shekhar11** 1 year, 9 months ago

Question in my exam 22sept 2023

upvoted 2 times

  **nikipediaa** 1 year, 9 months ago

me too today

upvoted 1 times

🗨️ 👤 **Tarajee** 1 year, 9 months ago

On my exam 2023sept  
upvoted 2 times

🗨️ 👤 **kayvg** 1 year, 10 months ago

1. Monitored container

> You want to track changes of the data you store

2. Lease container

> A lease container is responsible for maintaining information about which workers are processing which data from the Monitored Container

3. Host

> They represent the worker instances responsible for processing changes from the change feed

4. Delegates

> Methods you define to handle actual changes detected by the change feed (so when the change feed receives changes, it invokes delegate methods)

The difference between a monitored container and a host is that a host actually distributes work to delegates, while a monitored container tracks changes in data and gives that info to the change feed processor.

upvoted 5 times

🗨️ 👤 **KrishTeam** 1 year, 11 months ago

please provide correct answers

upvoted 1 times

🗨️ 👤 **unraval** 1 year, 11 months ago

The discussions are meant to help prepare for the exam so please try to be helpful

upvoted 2 times

🗨️ 👤 **unraval** 1 year, 11 months ago

People are commenting here to be funny, but can they care to provide an explanation?

upvoted 2 times

🗨️ 👤 **Chris2349** 2 years ago

Received this on 15th of June 2023. Went with the given answer.

upvoted 2 times

🗨️ 👤 **Videira** 2 years, 4 months ago

On my exam 2023-02-25

upvoted 3 times

🗨️ 👤 **mdg3501** 2 years, 6 months ago

got this on 2022-12-7

upvoted 3 times

🗨️ 👤 **TheExamMaster2020** 2 years, 7 months ago

Did my exam on 15th November 2022. This question was on it.

upvoted 5 times

🗨️ 👤 **BennyJai** 2 years, 6 months ago

I also Got same in 5 dec exam.

upvoted 2 times

🗨️ 👤 **herrmutig** 3 years, 2 months ago

Answer is correct.

\*Note that the "Host" Component should be called Compute Instance instead.

<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-processor>

upvoted 6 times

## HOTSPOT -

You are developing a web application that will use Azure Storage. Older data will be less frequently used than more recent data. You need to configure data storage for the application. You have the following requirements:

- ⇒ Retain copies of data for five years.
- ⇒ Minimize costs associated with storing data that is over one year old.
- ⇒ Implement Zone Redundant Storage for application data.

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Requirement	Solution
Configure an Azure Storage account	<div>▼</div> <ul style="list-style-type: none"> <li>Implement Blob Storage</li> <li>Implement Azure Cosmos DB</li> <li>Implement Storage (general purpose v1)</li> <li>Implement StorageV2 (general purpose v2)</li> </ul>
Configure data retention	<div>▼</div> <ul style="list-style-type: none"> <li>Snapshot blobs and move them to the archive tier</li> <li>Set a lifecycle management policy to move blobs to the cool tier</li> <li>Use AzCopy to copy the data to an on-premises device for backup</li> <li>Set a lifecycle management policy to move blobs to the archive tier</li> </ul>


### Suggested Answer:

#### Answer Area

Requirement	Solution
Configure an Azure Storage account	<div>▼</div> <ul style="list-style-type: none"> <li>Implement Blob Storage</li> <li>Implement Azure Cosmos DB</li> <li>Implement Storage (general purpose v1)</li> <li>Implement StorageV2 (general purpose v2)</li> </ul>
Configure data retention	<div>▼</div> <ul style="list-style-type: none"> <li>Snapshot blobs and move them to the archive tier</li> <li>Set a lifecycle management policy to move blobs to the cool tier</li> <li>Use AzCopy to copy the data to an on-premises device for backup</li> <li>Set a lifecycle management policy to move blobs to the archive tier</li> </ul>

#### Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers> <https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy?toc=/azure/storage/blobs/toc.json>

 **gmishra88** Highly Voted 2 years, 8 months ago

So, because it is ZRS, and that does not support archive tier, it cannot be moved to archive tier even though the questions mention the red-herring key-word "infrequently accessed" (which triggers feelings for archive tier). For no logically apparent reason Microsoft decided not to support archive tier in ZRS and unfortunately I have to remember that Microsoft "feature"?


upvoted 41 times

 **c75e123** 2 months ago

Only storage accounts that are configured for LRS, GRS, or RA-GRS support moving blobs to the archive tier. The archive tier isn't supported for ZRS, GZRS, or RA-GZRS accounts.



<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#archive-access-tier>

upvoted 1 times

  **Occ50bf** 10 months, 3 weeks ago

Archive data is physically disconnected from the web, if it has to be zone-redundant then it needs to be online, so it cannot be archived.

upvoted 1 times

  **OPT\_001122** 2 years, 7 months ago

ZRS, and that does not support archive tier - this is key point, Thanks

upvoted 13 times

  **surprise0011** 2 years, 2 months ago

also GZRS and RA-GZRS does not support archive tier

<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#archive-access-tier>

upvoted 3 times

  **Mick1z8** Highly Voted  1 year, 6 months ago

Correct answer. Got this in exam on 30/12/2023.

Case study: Contoso Ltd.

Total questions: 46

Time: 1:40 minutes

Score: 940

43 questions from Exam Topics. Just 3 questions outside of it.

upvoted 11 times

  **wafa\_chaari** Most Recent  5 months ago

correct.

The archive tier for Blob Storage isn't currently supported for ZRS, GZRS, or RA-GZRS accounts. Unmanaged disks don't support ZRS or GZRS.

General-purpose v2 storage accounts support the latest Azure Storage features and incorporate all of the functionality of general-purpose v1 and Blob storage accounts. General-purpose v2 accounts are recommended for most storage scenarios. General-purpose v2 accounts deliver the lowest per-gigabyte capacity prices for Azure Storage, as well as industry-competitive transaction prices. General-purpose v2 accounts support default account access tiers of hot or cool and blob level tiering between hot, cool, or archive.

upvoted 1 times

  **chrisbele** 7 months, 1 week ago

The answer is correct!

upvoted 1 times

  **neelkanths** 1 year, 2 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...



answer is correct...

upvoted 4 times

  **Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 2 times

  **applepie** 1 year, 11 months ago



got this question today, go with the provided answer - 7/30/2023, score 895/1000

upvoted 1 times

  **Priya0703** 2 years, 4 months ago

On exam 20-02-2023

upvoted 3 times

  **Esward** 2 years, 5 months ago

1. Implement storage V2

2. Set lifecycle management policy to move to cool tier

which means given answers are correct

upvoted 2 times

  **Hasti123** 1 year, 9 months ago

Why not Blob storage ?

upvoted 2 times

🗳️ 👤 **Woksi** 1 year, 4 months ago  
Blob Storage is not a type of \*Account\*  
upvoted 2 times

🗳️ 👤 **18Marks** 2 years, 6 months ago  
Got in exam. go with given answer  
upvoted 11 times

🗳️ 👤 **TonyMel** 2 years, 3 months ago  
correct, in 2023Mar24, score: 904/1000  
upvoted 3 times

🗳️ 👤 **adilkhan** 2 years, 3 months ago  
were all the questions from the exam topic?  
upvoted 3 times

🗳️ 👤 **Alluru** 2 years, 8 months ago  
Given answer is correct. Data retention policy relay on cool tier  
upvoted 4 times

🗳️ 👤 **gmishra88** 2 years, 8 months ago  
Microsoft plays on the unfortunate choice of words in their documentation: "Rarely used", "Infrequently used". The difference I cannot find  
upvoted 6 times

🗳️ 👤 **OPT\_001122** 2 years, 9 months ago  
Given answer is correct or not?  
upvoted 1 times

🗳️ 👤 **micro9000** 2 years, 7 months ago  
Yup, the answer is correct  
Based on this: <https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview?source=recommendations>  
Only storage accounts that are configured for LRS, GRS, or RA-GRS support moving blobs to the archive tier. The archive tier isn't supported for ZRS, GZRS, or RA-GZRS accounts. For more information about redundancy configurations for Azure Storage, see Azure Storage redundancy.  
upvoted 5 times

🗳️ 👤 **Enigma\_\_\_** 2 years, 9 months ago  
Only storage accounts that are configured for LRS, GRS, or RA-GRS support moving blobs to the Archive tier. The Archive tier isn't supported for ZRS, GZRS, or RA-GZRS accounts.  
<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview?source=recommendations>  
upvoted 3 times

🗳️ 👤 **finnishr** 2 years, 9 months ago  
Correct. Can't pick the archive option since ZRS needs to be used.  
upvoted 2 times

🗳️ 👤 **RochaG2** 2 years, 10 months ago  
Data in all tiers, including the Archive tier, is always copied from the primary to the secondary during geo-replication. The Archive tier for Blob Storage is currently supported for LRS, GRS, and RA-GRS accounts, but not for ZRS, GZRS, or RA-GZRS accounts.

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>  
upvoted 2 times

🗳️ 👤 **sghaha** 3 years, 2 months ago  
<https://docs.microsoft.com/ko-kr/azure/storage/common/storage-account-overview>  
upvoted 1 times

🗳️ 👤 **CelICS** 2 years, 5 months ago  
here is en-us link : <https://learn.microsoft.com/en-us/azure/storage/common/storage-account-overview>  
upvoted 1 times

## HOTSPOT -

A company develops a series of mobile games. All games use a single leaderboard service.

You have the following requirements:

- ⇒ Code must be scalable and allow for growth.
- ⇒ Each record must consist of a playerId, gameId, score, and time played.
- ⇒ When users reach a new high score, the system will save the new score using the SaveScore function below.

Each game is assigned an Id based on the series title.

▪

You plan to store customer information in Azure Cosmos DB. The following data already exists in the database:

PartitionKey	RowKey	Email
Harp	Walter	wharp@contoso.com
Smith	Steve	ssmith@contoso.com
Smith	Jeff	jsmith@contoso.com

You develop the following code to save scores in the data store. (Line numbers are included for reference only.)

```

01 public void SaveScore(string gameId, string playerId, int score, long timePlayed)
02 {
03     CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);
04     CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
05     CloudTable table = tableClient.GetTableReference("scoreTable");
06     table.CreateIfNotExists();
07     var scoreRecord = new PlayerScore(gameId, playerId, score, timePlayed);
08     TableOperation insertOperation = TableOperation.Insert(scoreRecord);
09     table.Execute(insertOperation);
10 }

```

You develop the following code to query the database. (Line numbers are included for reference only.)

```

01 CloudTableClient tableClient = account.CreateCloudTableClient();
02 CloudTable table = tableClient.GetTableReference("people");
03 TableQuery<CustomerEntity> query = new TableQuery<CustomerEntity>()
04     .Where(TableQuery.CombineFilters(
05         TableQuery.GenerateFilterCondition(PartitionKey, QueryComparisons.Equal, "Smith"),
06         TableOperators.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "ssmith@contoso.com")
07     ));
08 await table.ExecuteQuerySegmentedAsync<CustomerEntity>(query, null);

```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
SaveScore will work with Cosmos DB.	<input type="radio"/>	<input type="radio"/>
SaveScore will update and replace a record if one already exists with the same playerId and gameId.	<input type="radio"/>	<input type="radio"/>
Leader board data for the game will be automatically partitioned using gameId.	<input type="radio"/>	<input type="radio"/>
SaveScore will store the values for the gameId and playerId parameters in the database.	<input type="radio"/>	<input type="radio"/>

### Answer Area

	Statements	Yes	No
Suggested Answer:	SaveScore will work with Cosmos DB.	<input checked="" type="radio"/>	<input type="radio"/>
	SaveScore will update and replace a record if one already exists with the same playerId and gameld.	<input type="radio"/>	<input checked="" type="radio"/>
	Leader board data for the game will be automatically partitioned using gameld.	<input type="radio"/>	<input checked="" type="radio"/>
	SaveScore will store the values for the gameld and playerId parameters in the database.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: Yes -

Create a table.

A CloudTableClient object lets you get reference objects for tables and entities. The following code creates a CloudTableClient object and uses it to create a new

CloudTable object, which represents a table

// Retrieve storage account from connection-string.

CloudStorageAccount storageAccount =

CloudStorageAccount.parse(storageConnectionString);

// Create the table client.

CloudTableClient tableClient = storageAccount.createCloudTableClient();

// Create the table if it doesn't exist.

String tableName = "people";

CloudTable cloudTable = tableClient.getTableReference(tableName); cloudTable.createIfNotExists();

Box 2: No -

New records are inserted with TableOperation.insert. Old records are not updated.

To update old records TableOperation.insertOrReplace should be used instead.

Box 3: No -

Box 4: Yes -


Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-java>

  **yukkki** Highly Voted 1 year, 11 months ago

given answer is correct.

upvoted 16 times

  **[Removed]** 1 year, 5 months ago


Incorrect. First one is "No"

upvoted 4 times

  **Esward** 1 year, 11 months ago



yes, you are correct! given answers are correct

upvoted 3 times

  **OPT\_001122** Highly Voted 2 years, 2 months ago

the given answer seems correct

upvoted 8 times

  **vizay** Most Recent 2 months ago

Think of CloudStorageAccount as a key that can open the simpler "Azure Storage Account" toolbox and also the "table-mimicking" machine inside the more powerful "Azure Cosmos DB" workshop. However, it won't work on the other machines in the "Azure Cosmos DB" workshop (the ones for SQL, MongoDB, etc.). Those need their own special keys and instructions (different SDKs).

upvoted 1 times

  **Woksi** 10 months, 3 weeks ago

Is the implementation of the PlayerScore viewable in the exam? Wouldn't questions about partition/row key would be settled there?

upvoted 2 times

🗨️ 👤 **raymond\_abcd** 11 months, 1 week ago

See: <https://learn.microsoft.com/en-us/dotnet/api/overview/azure/data.tables-readme?view=azure-dotnet>

upvoted 3 times

🗨️ 👤 **raymond\_abcd** 11 months, 1 week ago

CloudStorageClient is deprecated for CosmosDb it is now changed to TableServiceClient and then you have TableClient to store the data. So probably this question will be different in the upcoming exam

upvoted 3 times

🗨️ 👤 **Azr0112** 1 year, 3 months ago

Yes, SaveScore will work with CosmosDb.

This is an example of Azure Cosmos DB for Table.

<https://learn.microsoft.com/en-us/azure/cosmos-db/table/quickstart-dotnet?tabs=azure-cli%2Cwindows>

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 5 months ago

No, No, No, Yes

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 11 months ago

Inserting will fail, this would require InsertOrReplace to work

upvoted 3 times

🗨️ 👤 **bbq598** 2 years ago

For box 2 there will be an exception if the same partition key and row key already exist.

upvoted 2 times

🗨️ 👤 **warchoon** 1 year, 10 months ago

So you need to answer box 3 first ;)

upvoted 1 times

🗨️ 👤 **Yumi21** 2 years, 1 month ago

It seems that in the code no partition key for scoreTable is specified, which means that here we are using a single-partition collection. I would say in this case 3rd option is No, so no automatical partitioning will happen.

upvoted 2 times

🗨️ 👤 **TheExamMaster2020** 2 years, 2 months ago

Y,N,Y,Y

upvoted 1 times

🗨️ 👤 **coffecold** 2 years, 2 months ago

What code should be there if doing automatic partitioning (C) remains a total secret to me. Should be somewhere on the table level...

upvoted 4 times

🗨️ 👤 **DivyaRajkumar** 2 years, 3 months ago

The given answer looks right to me,

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table?view=azure-dotnet>

upvoted 2 times

🗨️ 👤 **azurepaul** 2 years, 3 months ago

CloudTableClient is for Table storage - CosmosClient is for Cosmos DB I don't think it will work with Cosmos

upvoted 6 times

🗨️ 👤 **PrepX** 1 year, 6 months ago

A = No

SaveScore() method uses a storage account connectionString:

CloudStorageAccount storageAccount = CloudStorageAccount.Parse(connectionString);

and uses storageAccount to create CloudTableClient.

This is not a CosmosDB connection string, CosmosDB requires a Uri + credential (not given in given code!) So given code will not work 100% to connect to CosmosDB even though it uses a compatible API.

upvoted 3 times



🗨️ 👤 **Enigma\_\_** 2 years, 3 months ago

Please do a simple bing search and you'll see it is Cosmos Db

<https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table.cloudtableclient?view=azure-dotnet>

upvoted 1 times

🗨️ 👤 **ArturKon** 2 years, 3 months ago

Cosmos DB offers Table API, you can use CosmosClient to connect to this.

upvoted 2 times

🗨️ 👤 **ArturKon** 2 years, 3 months ago

<https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table.cloudtableclient?view=azure-dotnet>

upvoted 1 times

🗨️ 👤 **finnishr** 2 years, 3 months ago

Correct 100%!

upvoted 2 times

🗨️ 👤 **BogdanG** 2 years, 8 months ago

Answer is Correct.

upvoted 3 times

You develop and deploy a web application to Azure App Service. The application accesses data stored in an Azure Storage account. The account contains several containers with several blobs with large amounts of data. You deploy all Azure resources to a single region. You need to move the Azure Storage account to the new region. You must copy all data to the new region. What should you do first?

- A. Export the Azure Storage account Azure Resource Manager template
- B. Initiate a storage account failover
- C. Configure object replication for all blobs
- D. Use the AzCopy command line tool
- E. Create a new Azure Storage account in the current region
- F. Create a new subscription in the current region

#### Suggested Answer: A

To move a storage account, create a copy of your storage account in another region. Then, move your data to that account by using AzCopy, or another tool of your choice and finally, delete the resources in the source region.

To get started, export, and then modify a Resource Manager template.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-account-move?tabs=azure-portal>

Community vote distribution

A (100%)

🗳️ 👤 **[Removed]** Highly Voted 1 year, 8 months ago

Selected Answer: A

We can create a new Storage account in the new region, using the existing storage account ARM template. All we need to do is change the region name after exporting the ARM of existing account..

upvoted 13 times

🗳️ 👤 **gmishra88** 1 year, 2 months ago

And the name of the storage account which needs to be unique. I think all these questions are about some sentence in some azure documentation.

upvoted 4 times

🗳️ 👤 **Nhiendo** Highly Voted 1 year ago

Step One is always "Export"

upvoted 10 times

🗳️ 👤 **vizay** Most Recent 2 months ago

Selected Answer: A

export the Azure Storage Account's ARM (Azure Resource Manager) template – basically, the "blueprint" of your storage account that you can reuse or redeploy.

upvoted 1 times

🗳️ 👤 **OPT\_001122** 1 year, 1 month ago

Selected Answer: A

A. Export the Azure Storage account Azure Resource Manager template

upvoted 3 times

🗳️ 👤 **TheExamMaster2020** 1 year, 1 month ago

Did my exam on 15th November 2022. This question was on it.

upvoted 6 times

🗳️ 👤 **BogdanG** 1 year, 8 months ago

Answer is correct.

upvoted 10 times

## HOTSPOT -

You are developing an application to collect the following telemetry data for delivery drivers: first name, last name, package count, item id, and current location coordinates. The app will store the data in Azure Cosmos DB.

You need to configure Azure Cosmos DB to query the data.

Which values should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

### Configuration Parameter

### Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

## Answer Area

### Configuration Parameter

### Value

Azure Cosmos DB API

	▼
Gremlin	
Table API	
Core (SQL)	

Suggested Answer:

Azure Cosmos DB partition key

	▼
first name	
last name	
package count	
item id	

Box 1: Core (SQL)



Core(SQL) API stores data in document format. It offers the best end-to-end experience as we have full control over the interface, service, and

Box 2: item id -  
item id is a unique identifier and is suitable for the partition key.  
Reference:  
<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>  
<https://docs.microsoft.com/en-us/azure/cosmos-db/partitioning-overview>

Without knowing the functionality or the usage pattern or what it is for. Good lord, Microsoft upvoted 43 times

Got this in 09/22 , went with SQL and Item Id, score 927.  
upvoted 24 times

Got this today.  
Went with answer here.  
Score 927  
upvoted 7 times

  **Gauravbio** 1 year, 1 month ago

Core sql api removed now, it should be no sql api



upvoted 1 times

On my exam 2023sept  
upvoted 8 times


Had this question in today's exam: 2023-07-26  
upvoted 5 times



Answer seems correct, but the question is very bad. It doesn't even tell about the usage, so it could be Table API as well  
upvoted 7 times

Got this question in the exam on 16/03/2023. Went with SQL and Item Id . Make sure to prepare for case studies. I got city and lights case study.  
upvoted 5 times

  **MarcoStewart** 2 years ago  
where can one find these case studies?  
upvoted 3 times

Page 32 and above  
upvoted 1 times

 **Puja079888** 1 year, 4 months ago  
did contributor access is needed to pass the exam?  
upvoted 1 times

  **Christian\_garcia\_martin** 10 months, 2 weeks ago  
yes it is  
upvoted 1 times

They changed names again. "Core (SQL)" is "Api for NoSQL" now. Its wonderful to choose "SQL" for "NoSQL".  
<https://www.c-sharpcorner.com/article/road-to-az-2044/#:~:text=Core%20SQL%20API%2C%20default%20API%20for%20using%20Azure%20Cosmos%20DB%20enables%20querying%20your%20data%20with%20SQL,https://learn.microsoft.com/en-us/azure/cosmos-db/choose-api#coresql-api#:~:text=API%20for%20NoSQL%20is%20native%20to%20Azure%20Cosmos%20DB>  
 upvoted 12 times

🗨️ 👤 **Esward** 2 years, 5 months ago

given answers are correct

upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

SQL and Item Id

upvoted 2 times

🗨️ 👤 **rol204** 2 years, 8 months ago

correct

upvoted 2 times

🗨️ 👤 **kampatra** 2 years, 9 months ago

Answer is correct:

<https://docs.microsoft.com/en-us/azure/cosmos-db/choose-api>

upvoted 4 times

🗨️ 👤 **dtctx** 3 years, 2 months ago

Received this in test on 4/29 and passed the test.

Went with Table API and Item ID. I do not know if Table API is correct, but I am confident that Item ID is.

upvoted 4 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 7 months ago

Table API is for manage tables, SQL is for queries

upvoted 1 times

🗨️ 👤 **sghaha** 3 years, 2 months ago

<https://docs.microsoft.com/ko-kr/azure/cosmos-db/choose-api>

upvoted 1 times

## DRAG DROP -

You are implementing an Azure solution that uses Azure Cosmos DB and the latest Azure Cosmos DB SDK. You add a change feed processor to a new container instance.

You attempt to read a batch of 100 documents. The process fails when reading one of the documents. The solution must monitor the progress of the change feed processor instance on the new container as the change feed is read. You must prevent the change feed processor from retrying the entire batch when one document cannot be read.

You need to implement the change feed processor to read the documents.

Which features should you use? To answer, drag the appropriate features to the cored requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each cored selection is worth one point.

Select and Place:

## Features

Change feed estimator

Dead-letter queue

Deployment unit

Lease container

## Answer Area

## Requirement

Monitor the progress of the change feed processor

Prevent the change feed processor from retrying the entire batch when one document cannot be read

## Feature

## Suggested Answer:

## Features

Change feed estimator

Dead-letter queue

Deployment unit

Lease container

## Answer Area

## Requirement

Monitor the progress of the change feed processor

Prevent the change feed processor from retrying the entire batch when one document cannot be read

## Feature

Change feed estimator

Dead-letter queue

Box 1: Change feed estimator -

You can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed or use the life cycle notifications to detect underlying failures.

Box 2: Dead-letter queue -

To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to a dead-letter queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The dead-letter queue might be another Cosmos container. The exact data store does not matter, simply that the unprocessed changes are persisted.


Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-processor>

 **finnishr**  2 years, 10 months ago

The answer is correct.

upvoted 18 times

 **ns4098** 9 months, 1 week ago

1. Monitor the progress of the change feed processor

Feature: Lease container

Reason: The Lease container is responsible for tracking the progress of the change feed processor. It maintains state information, ensuring that the change feed processor knows where it left off and can resume from the correct point in the feed. This allows the system to monitor the progress and distribute the workload across multiple instances if needed.

2. Prevent the change feed processor from retrying the entire batch when one document cannot be read

Feature: Dead-letter queue

Reason: The Dead-letter queue allows you to handle documents that cannot be processed. If a document cannot be read or processed, instead of failing the entire batch, the problematic document is moved to a dead-letter queue, allowing the change feed processor to continue processing the rest of the batch without retrying the failed document.

upvoted 1 times

🗲️ 👤 **sarmaria** Highly Voted 👍 2 years, 3 months ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 13 times

🗲️ 👤 **Stann07** Most Recent 🕒 1 year, 6 months ago

On my exam Dec 18. went with the given answer. scored 842

upvoted 3 times

🗲️ 👤 **AndySmith** 1 year, 7 months ago

On exam 3-Nov-2023. Went with proposed answer - 932/1000.

1) Change feed estimator

2) Dead letter queue

upvoted 4 times

🗲️ 👤 **dddddd111** 1 year, 9 months ago

I got this same question. Provided answers are correct. (Note: I failed the exam 20/9/23. I only scored 644 and I felt bad. I think because many questions here in Examtopics are not accurate. I suggest following the most voted answers and don't just not rely on Examtopics answers. At the beginning of the exam, you will be asked which programming languages you want to use. C#/Python. I chose C#. Also, I just want to add that some questions here are really in the actual exams, but the choices are written and formatted differently. Please be aware of that. Goodluck. I feel bad for failing it, but I want to retake next month. I will try Python. T\_T

upvoted 10 times

🗲️ 👤 **Nitin23** 1 year, 9 months ago

this is a bad site

upvoted 2 times

🗲️ 👤 **kayvg** 1 year, 10 months ago

The answer is correct

A change feed estimator is used to monitor the progress of your change feed processor instances as they read the change feed

A dead-letter queue is holding a queue for messages that cannot be delivered to their destination

A deployment unit is to provide a container for an application or service, which is not relevant

The lease container is used to coordinate processing the change feed, which is not relevant as well

upvoted 1 times

🗲️ 👤 **kvtborad** 1 year, 10 months ago

I got this question on 6th August 2023. chose highly voted. passed with 904. I got Case study: city and Lights. All questions are from ExamTopics.

upvoted 2 times

🗲️ 👤 **JH81** 2 years ago

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 2 times

🗲️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

My case:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 2 times

🗨️ 👤 **narenazure** 2 years, 5 months ago

The answer is correct.

To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The errored-message queue might be another Azure Cosmos DB container. The exact data store does not matter, simply that the unprocessed changes are persisted.

Error handling section - <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet>

upvoted 1 times

🗨️ 👤 **Esward** 2 years, 5 months ago

Given answers are correct!

upvoted 1 times

🗨️ 👤 **Esward** 2 years, 5 months ago

To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The errored-message queue might be another Azure Cosmos DB container. The exact data store does not matter, simply that the unprocessed changes are persisted.

In addition, you can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed or use the life cycle notifications to detect underlying failures.

upvoted 2 times

🗨️ 👤 **Sreedharc** 2 years, 7 months ago

got in 11/11/2022

upvoted 3 times

🗨️ 👤 **vcfvct** 2 years, 8 months ago

Got this on 10/21/2022

upvoted 3 times

🗨️ 👤 **OPT\_001122** 2 years, 8 months ago

thanks for mentioning the date

upvoted 1 times

🗨️ 👤 **gmishra88** 2 years, 8 months ago

Sure, this is also taken from a certification. But where does a dead letter queue come in change feed processor. If this comes, yes, this is the correct answer. But it all doesn't make any sense

upvoted 1 times

🗨️ 👤 **coffecold** 2 years, 8 months ago

It is do-it-yourself stuff:

"To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to a dead-letter queue "

<https://learn.microsoft.com/en-us/azure/cosmos-db/sql/change-feed-processor?tabs=dotnet>

upvoted 3 times

🗨️ 👤 **Pentagon** 2 years, 8 months ago

If just for saving the last position, I think it should be lease container.

"When the delegate finishes processing the changes successfully, update the lease store with the latest processed point in time and go to #1"

Dead-letter queue seems to be prevent re-trying too much instead of resuming at the last done position.

upvoted 1 times



## HOTSPOT -

You are developing an application that uses a premium block blob storage account. The application will process a large volume of transactions daily. You enable

Blob storage versioning.

You are optimizing costs by automating Azure Blob Storage access tiers. You apply the following policy rules to the storage account. (Line numbers are included for reference only.)

```

01 {
02   "rules" : [
03     {
04       "name" : "versionRule",
05       "enabled" : true,
06       "type" : "Lifecycle",
07       "definition" : {
08         "actions" : {
09           "version" : {
10             "tierToCool" : {
11               "daysAfterCreationGreaterThan" : 60
12             },
13             "delete" : {
14               "daysAfterCreationGraterThan" : 365
15             }
16           }
17         },
18         "filters" : {
19           "blobTypes" : [ "blockBlob" ], "prefixMatch" : [ "transactions" ]
20         }
21       }
22     }
23   ]
24 }
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

Statements	Yes	No
Block blobs prefixed with <b>transactions</b> will transition blobs that have not been modified in over 60 days to cool storage, and delete blobs not modified in 365 days	<input type="radio"/>	<input type="radio"/>
Blobs are moved to cool storage if they have not been accessed for 60 days	<input type="radio"/>	<input type="radio"/>
The policy rule tiers previous versions within a container named transactions that are 60 days or older to the cool tier and deletes previous versions that are 365 days or older	<input type="radio"/>	<input type="radio"/>
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

Answer Area

Statements	Yes	No
Block blobs prefixed with <b>transactions</b> will transition blobs that have not been modified in over 60 days to cool storage, and delete blobs not modified in 365 days	<input type="radio"/>	<input checked="" type="radio"/>
Blobs are moved to cool storage if they have not been accessed for 60 days	<input type="radio"/>	<input checked="" type="radio"/>
The policy rule tiers previous versions within a container named transactions that are 60 days or older to the cool tier and deletes previous versions that are 365 days or older	<input checked="" type="radio"/>	<input type="radio"/>
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: No -

Would be true if daysAfterModificationGreaterThan was used, but here daysAfterCreationGreaterThan

Box 2: No -

Would need to use the daysAfterLastAccessTimeGreaterThan predicate.

Box 3: Yes -

Box 4: Yes -

With the lifecycle management policy, you can:

Transition blobs from cool to hot immediately when they are accessed, to optimize for performance.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

 **baroo1**  2 years, 4 months ago

No - Not modified, created.

No - Not accessed, created.

Yes - Rules are matching the statement. The prefix "transactions" can be applicable for containers as well. "container" / "container/blob" or "blob" can be used under this context.

Source: <https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#archive-data-after-ingest>

No - "enableAutoTierToHotFromCool": "true" should be enabled.

upvoted 29 times

 **Dani\_ac7**  2 years, 10 months ago

With this image, all answers are NO:

- Container named transaction is not in code

- is no present line "enableAutoTierToHotFromCool": true

upvoted 26 times

 **Tenk14** 2 years, 9 months ago

I guess, third statement (The policy rule tiers..) result is Yes.

Container name "transactions" is in prefixMatch.

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview#archive-data-after-ingest>

Solution is:

- No

- No

- Yes

- No

upvoted 63 times

 **ChiragShah4885** 2 years, 9 months ago

Container name "transactions" is in prefixMatch means its name must start from "transactions". Its name is not "transactions". Dani\_ac7's answer looks correct  
upvoted 2 times

  **ArturKon** 2 years, 9 months ago

Looks like 3rd is NO, according to filter's guide:

"Filter blobs by name or first letters. To find items in a specific container, enter the name of the container followed by a forward slash, then the blob name or first letters. For example, to show all blobs starting with "a", type: "mycontainer/a"."



So it looks like we are searching blobs with "transactions" prefix in all containers in 3rd.

upvoted 1 times

  **cwn53066** 2 years, 9 months ago

But name "transactions" satisfies rule `name must start from "transactions"`, why it doesn't fit?

upvoted 5 times

  **Esward** 2 years, 5 months ago

Correct: if you want to match the blobs within a specific container, you should mention the container name/blob name

```
{
  "rules": [
    {
      "name": "agingRule",
      "enabled": true,
      "type": "Lifecycle",
      "definition": {
        "filters": {
          "blobTypes": [ "blockBlob" ],
          "prefixMatch": [ "sample-container/blob1", "container2/blob2" ]
        },
        "actions": {
          "baseBlob": {
            "tierToCool": { "daysAfterModificationGreaterThan": 30 },
            "tierToArchive": { "daysAfterModificationGreaterThan": 90 }
          }
        }
      }
    }
  ]
}
```

upvoted 1 times

  **Hendrikdb** 1 year, 4 months ago

wrong:

An array of strings for prefixes to be matched. Each rule can define up to 10 case-sensitive prefixes. A prefix string must start with a container name. For example, if you want to match all blobs under <https://myaccount.blob.core.windows.net/sample-container/blob1/...> for a rule, the prefixMatch is sample-container/blob1

upvoted 1 times

  **DamuKeesh**  5 months ago

The solution is

-No

-No

-Yes

-No

upvoted 1 times

  **4bd3116** 10 months, 1 week ago

Block blobs prefixed with "transactions" will transition blobs that have not been modified in over 60 days to cool storage, and delete blobs not modified in 365 days.

Yes. This is correct based on the rule provided.

Blobs are moved to cool storage if they have not been accessed for 60 days.

No. The rule is based on the creation date, not access date.

The policy rule tiers previous versions within a container named "transactions" that are 60 days or older to the cool tier and deletes previous versions that are 365 days or older.

No. The rule applies to block blobs with the prefix "transactions," not specifically to previous versions within a container.

Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.

No. As you correctly pointed out, the rule does not specify "enableAutoTierToHotFromCool" as true.

So, the correct answers should be: Yes, No, No, No.

upvoted 1 times

🗳️ 👤 **Christian\_garcia\_martin** 10 months, 2 weeks ago

NO , NO ,YES ,NO . Last one is NO because you don't see "enableAutoTierToHotFromColl" : true

upvoted 3 times

🗳️ 👤 **jobolesonihal** 1 year, 3 months ago

policy doesn't mention baseblob.

Solution is:

- No

- No

- Yes

- No

upvoted 4 times

🗳️ 👤 **SSR999** 1 year, 4 months ago

Premium Block Blobs wont support access tiers and Lifecycle management policies (tiering)

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-feature-support-in-storage-accounts#premium-block-blob-accounts>

upvoted 6 times

🗳️ 👤 **NPE\_** 1 year, 8 months ago

3rd is YES, check the reference:

<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

On the page it says: "A prefix string must start with a container name. For example, if you want to match all blobs under

<https://myaccount.blob.core.windows.net/sample-container/blob1/...> for a rule, the prefixMatch is sample-container/blob1." This means that the container name is "transactions"

So the only one that make me confusing is the 4th. According to the discussion, you have to manually enable "enableAutoTierToHotFromCool". And it is also mentioned above in the link that "The enableAutoTierToHotFromCool action is available only when used with the daysAfterLastAccessTimeGreaterThan run condition. " As in our case, the prerequisite for the auto tier to hot from cool is not even fulfilled, so it is not possible to expect an "AutoTierToHotFromCool" effect.

I would go with a NO for the 4th.

upvoted 3 times

🗳️ 👤 **nikipediaa** 1 year, 9 months ago

Question was on exam 2023-09-26

upvoted 1 times

🗳️ 👤 **nikipediaa** 1 year, 9 months ago

sorry, not this one, but similar with containers and numbers :)

upvoted 1 times

🗳️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

Without 4th question.

my cases also:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 2 times

🗉 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

My case:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 1 times

🗉 👤 **surprise0011** 2 years, 2 months ago

received 2023-04-17 went N,N,Y, score 926

last box was not there only first three

upvoted 8 times

🗉 👤 **surprise0011** 2 years, 2 months ago

premium block blob storage does not support access tiers. This is confusing

upvoted 4 times

🗉 👤 **Saluk\_DE** 2 years, 3 months ago

Question was in Exam 2023-03-30

upvoted 1 times

🗉 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 4 times

🗉 👤 **rasojol523** 2 years, 3 months ago

Box 3: No ?

'DayAfterCreationGraterThan 60' does not include 'just 60 days'. Therefore, if something has been made for 60 days, it would not be included in '60 days or later', so box 3 would be NO ?

upvoted 2 times

🗉 👤 **Ayman99** 2 years, 4 months ago

Was on the exam 27.02.2023

Went with NNYN.

Score 870

upvoted 3 times

An organization deploys Azure Cosmos DB.

You need to ensure that the index is updated as items are created, updated, or deleted.

What should you do?

- A. Set the indexing mode to Lazy.
- B. Set the value of the automatic property of the indexing policy to False.
- C. Set the value of the EnableScanInQuery option to True.
- D. Set the indexing mode to Consistent.

**Suggested Answer: D**

Azure Cosmos DB supports two indexing modes:

Consistent: The index is updated synchronously as you create, update or delete items. This means that the consistency of your read queries will be the consistency configured for the account.

None: Indexing is disabled on the container.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/index-policy>

Community vote distribution

D (100%)

 **finnishr** Highly Voted 1 year, 10 months ago

Correct !!!

upvoted 14 times

 **kayvg** Highly Voted 10 months, 2 weeks ago

Answer: D (consistent)

Setting indexing mode to consistent ensures that index is updated in sync with data

Setting it to lazy means index is updated async, meaning you may not get latest version of data

Setting the value of the automatic property to false means you have to manually manage indexing of data

Setting EnableScanInQuery to true allows query operations to perform scans across all documents, this is not really relevant + it provides poor performance

upvoted 7 times


 **lyggwtwtczxnhznebw** Most Recent 11 months, 3 weeks ago

**Selected Answer: D**

Setting the indexing mode to Consistent ensures that the index is always kept up to date with the changes made to the items in Azure Cosmos DB.

With this mode, whenever an item is created, updated, or deleted, the index is automatically updated in real-time to reflect those changes.

upvoted 1 times

 **[Removed]** 1 year, 3 months ago

**Selected Answer: D**

D indexing mode ensures that the index is always up-to-date with the data stored in the container and changes are immediately reflected in the index.

upvoted 1 times

 **Esward** 1 year, 5 months ago

Given answer is correct.

Consistent: The index is updated synchronously as you create, update or delete items. This means that the consistency of your read queries will be the consistency configured for the account.

None: Indexing is disabled on the container. This mode is commonly used when a container is used as a pure key-value store without the need for secondary indexes. It can also be used to improve the performance of bulk operations. After the bulk operations are complete, the index mode can be set to Consistent and then monitored using the IndexTransformationProgress until complete.

upvoted 4 times

  **OPT\_001122** 1 year, 7 months ago

**Selected Answer: D**

D. Set the indexing mode to Consistent.

upvoted 3 times

You are developing a .Net web application that stores data in Azure Cosmos DB. The application must use the Core API and allow millions of reads and writes.

The Azure Cosmos DB account has been created with multiple write regions enabled. The application has been deployed to the East US2 and Central US regions.

You need to update the application to support multi-region writes.

What are two possible ways to achieve this goal? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Update the ConnectionPolicy class for the Cosmos client and populate the PreferredLocations property based on the geo-proximity of the application.
- B. Update Azure Cosmos DB to use the Strong consistency level. Add indexed properties to the container to indicate region.
- C. Update the ConnectionPolicy class for the Cosmos client and set the UseMultipleWriteLocations property to true.
- D. Create and deploy a custom conflict resolution policy.
- E. Update Azure Cosmos DB to use the Session consistency level. Send the SessionToken property value from the FeedResponse object of the write action to the end-user by using a cookie.

**Suggested Answer: CD**

C: The UseMultipleWriteLocations of the ConnectionPolicy class gets or sets the flag to enable writes on any locations (regions) for geo-replicated database accounts in the Azure Cosmos DB service.

Note: Once an account has been created with multiple write regions enabled, you must make two changes in your application to the ConnectionPolicy for the

Cosmos client to enable the multi-region writes in Azure Cosmos DB. Within the ConnectionPolicy, set UseMultipleWriteLocations to true and pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property based on the geo-proximity from location passed in. If a new region is later added to the account, the application does not have to be updated or redeployed, it will automatically detect the closer region and will auto-home on to it should a regional event occur.

Azure core API application "ConnectionPolicy class" cosmos db multiple write regions enabled


D: With multi-region writes, when multiple clients write to the same item, conflicts may occur. When a conflict occurs, you can resolve the conflict by using different conflict resolution policies.

Note: Conflict resolution policy can only be specified at container creation time and cannot be modified after container creation.

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.connectionpolicy> <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master> <https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-manage-conflicts>

Community vote distribution



**coffecold** Highly Voted 2 years, 8 months ago

The goal is

"You need to update the application to support multi-region writes",

that is enable multi-region writes (bool, option C) and add the regions (option A)

Then you have to apply the Conflict resolution policies. This can be LLW(default, not mentioned) or custom (option D).

Hence : there is only ONE way to to support multi-region writes (both apply C AND A) and there are subsequently TWO ways to apply the Conflict resolution policies (@ SQL) to solve write, update and delete conflicts of which one is mentioned in the question (D).

To support multi-region writes I would answer A and C , but they have to be set both, not one or the other.

See <https://learn.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master?tabs=api-async> and <https://learn.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies>

upvoted 20 times

**dgcc97** 2 years, 6 months ago

From the documentation you provided:

Within the ConnectionPolicy, set UseMultipleWriteLocations to true and pass the name of the region where the application is deployed to



ApplicationRegion. This will populate the PreferredLocations property based on the geo-proximity from location passed in. If a new region is later added to the account, the application does not have to be updated or redeployed, it will automatically detect the closer region and will auto-home on to it should a regional event occur.

I take that paragraph to mean that it automatically updates the PreferredLocations property, when you set the UseMultipleWriteLocations property with true and pass the region name.

That means you actually only need to update UseMultipleWriteLocations, thus C.

And as you mentioned, apply the Conflict resolution policies (D).

upvoted 4 times

  **warchoon** 2 years, 4 months ago

It's a trick. "you must make two changes in your application to the ConnectionPolicy" "pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property". I think this is an "update"

[https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-](https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async#:~:text=you%20must%20make%20two%20changes%20in%20your%20application%20to%20the%20ConnectionPolicy)

[async#:~:text=you%20must%20make%20two%20changes%20in%20your%20application%20to%20the%20ConnectionPolicy](https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async#:~:text=you%20must%20make%20two%20changes%20in%20your%20application%20to%20the%20ConnectionPolicy)

upvoted 2 times

  **baroo1**  2 years, 4 months ago

**Selected Answer: AC**

A + C =

Clearly stated in documentation that both are required: <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async>.

And according to the question "Each correct answer presents part of the solution". "Part"

Not D -> "Custom conflict resolution policy is available only for API for NoSQL accounts and can be set only at creation time." The question states the account was already created.

Source: <https://learn.microsoft.com/en-us/azure/cosmos-db/conflict-resolution-policies>

upvoted 13 times

  **8ac3742**  10 months, 4 weeks ago

D is not answer, Conflict resolution policy can only be specified at container creation time and cannot be modified after container creation, however the question mentions that "You need to update the application to support multi-region writes.", so we suppose that the container has been created before we update ConnectionPolicy, so D is not applicable.

The answer is C+A, set UseMultipleWriteLocation and PreferredLocation.

upvoted 1 times


  **singhdamodar3** 1 year, 6 months ago

**Selected Answer: CD**

C. Setting the UseMultipleWriteLocations property to true allows the Cosmos DB SDK to automatically route write operations to the closest available write region, improving write latency. This is a key configuration for enabling multi-region writes.

D - is relevant in scenarios where conflict resolution needs customization,

upvoted 1 times

  **applepie** 1 year, 11 months ago

got this question today, answer A,C - 7/30/2023, score 895/1000

upvoted 1 times

  **JH81** 2 years ago

**Selected Answer: AC**

Got this on 6/28/2023 and passed with 850. Went with A, C.

upvoted 2 times

  **deathRac3** 2 years, 1 month ago

**Selected Answer: AC**

A&C to support multi-region writes

upvoted 2 times

  **nekkilodeon** 2 years, 2 months ago

**Selected Answer: AC**

A & C

<https://learn.microsoft.com/en-us/azure/cosmos-db/sql/how-to-multi-master?tabs=api-async>

upvoted 2 times

🗳️ 👤 **adilkhan** 2 years, 2 months ago

A,C chat gpt

upvoted 2 times

🗳️ 👤 **vikas4239** 2 years, 4 months ago

A&C

ConnectionPolicy, set UseMultipleWriteLocations to true and pass the name of the region where the application is deployed to ApplicationRegion. This will populate the PreferredLocations property based on the geo-proximity from location passed in

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async>

upvoted 1 times

🗳️ 👤 **AkhiRk** 2 years, 4 months ago

Answer should be CD

The questions mentions .NET web application and if you check below link you can see for .net the option C goes

And for multi write scenarios conflict might arise and hence option D

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async>

upvoted 2 times

🗳️ 👤 **Esward** 2 years, 5 months ago

AC are the correct answers as per microsoft docs

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-multi-master?tabs=api-async>

upvoted 1 times

🗳️ 👤 **imanonion** 2 years, 5 months ago

**Selected Answer: AC**

Option C: Update the ConnectionPolicy class for the Cosmos client and set the UseMultipleWriteLocations property to true. This will enable the application to use multiple write regions when storing data in Azure Cosmos DB.

Option A: Update the ConnectionPolicy class for the Cosmos client and populate the PreferredLocations property based on the geo-proximity of the application. This will allow the application to use the closest write region to the application's location, improving performance and reducing latency. Other options, such as updating the consistency level or creating a custom conflict resolution policy, are not directly related to enabling multi-region writes and are not necessary for this goal.

upvoted 4 times

🗳️ 👤 **rotimislav** 2 years, 6 months ago

**Selected Answer: AC**

As stated by others

upvoted 2 times

🗳️ 👤 **Siepie** 2 years, 6 months ago

Correct answers should be A and C

ConnectionPolicy policy = new ConnectionPolicy

```
{  
    ConnectionMode = ConnectionMode.Direct,  
    ConnectionProtocol = Protocol.Tcp,  
    UseMultipleWriteLocations = true ----> C  
};
```

policy.SetCurrentLocation("West US 2"); ----> A

upvoted 4 times

🗳️ 👤 **Siepie** 2 years, 6 months ago

Sorry SetCurrentLocation is only the region in which the application is being deployed and where Azure Cosmos DB is replicated. But the ConnectionPolicy class does also have a property PreferredLocations: Gets and sets the preferred locations (regions) for geo-replicated database accounts in the Azure Cosmos DB service. For example, "East US" as the preferred location.

upvoted 1 times

🗳️ 👤 **OPT\_001122** 2 years, 8 months ago

C and D are the correct answers?

upvoted 1 times

🗨️ 👤 **P000JAAAAAAAAA** 2 years, 6 months ago

IMO C and D

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 9 months ago

The documentation says: " To enable multi-region writes in your application, set ApplicationRegion to the region in which the application is being deployed and where Cosmos DB is replicated"

Now the option to use that property is in .Net SDK V2. Microsoft says I do not need to know .Net for this exam and they ask questions from V2? So, I should try this in both v2 and v3 to be able to pass this exam.

upvoted 2 times

**HOTSPOT -**

You are developing a solution to store documents in Azure Blob storage. Customers upload documents to multiple containers. Documents consist of PDF, CSV,

Microsoft Office format and plain text files.

The solution must process millions of documents across hundreds of containers. The solution must meet the following requirements:

- ⇒ Documents must be categorized by a customer identifier as they are uploaded to the storage account.
- ⇒ Allow filtering by the customer identifier.
- ⇒ Allow searching of information contained within a document
- ⇒ Minimize costs.

You create and configure a standard general-purpose v2 storage account to support the solution.

You need to implement the solution.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area****Requirement****Solution**

Search and filter by customer identifier

	▼
Azure Cognitive Search	
Azure Blob index tags	
Azure Blob inventory policy	
Azure Blob metadata	

Search information inside documents

	▼
Azure Cognitive Search	
Azure Blob index tags	
Azure Blob inventory policy	
Azure Blob metadata	

**Suggested Answer:****Answer Area****Requirement****Solution**

Search and filter by customer identifier

	▼
Azure Cognitive Search	
Azure Blob index tags	
Azure Blob inventory policy	
Azure Blob metadata	

Search information inside documents

	▼
Azure Cognitive Search	
Azure Blob index tags	
Azure Blob inventory policy	
Azure Blob metadata	

As datasets get larger, finding a specific object in a sea of data can be difficult. Blob index tags provide data management and discovery capabilities by using key- value index tag attributes. You can categorize and find objects within a single container or across all containers in your storage account. As data requirements change, objects can be dynamically categorized by updating their index tags. Objects can remain in-place with their current container organization.


Box 2: Azure Cognitive Search -

Only index tags are automatically indexed and made searchable by the native Blob Storage service. Metadata can't be natively indexed or searched. You must use a separate service such as Azure Search.

Azure Cognitive Search is the only cloud search service with built-in AI capabilities that enrich all types of information to help you identify and explore relevant content at scale. Use cognitive skills for vision, language, and speech, or use custom machine learning models to uncover insights from all types of content.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-manage-find-blobs> <https://azure.microsoft.com/en-us/services/search/>

 **jeanfmc** Highly Voted 2 years, 10 months ago  
right

Azure Blob Index tags: <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-index-how-to?tabs=azure-portal>


Azure Cognitive Search: Search inside documents

upvoted 21 times

 **TheExamMaster2020** Highly Voted 2 years, 7 months ago

Did my exam on 15th November 2022. This question was on it.

upvoted 8 times

 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 2 times

 **ns4098** 9 months, 1 week ago

1. Search and filter by customer identifier:

Solution: Azure Blob index tags

Reason: Azure Blob index tags allow you to tag blobs with metadata such as a customer identifier. These tags make it easy to search and filter blobs based on specific attributes without the need to load or examine the entire content of the blob. This solution is efficient for filtering blobs by customer identifier.

2. Search information inside documents:

Solution: Azure Cognitive Search

Reason: Azure Cognitive Search is the appropriate solution for searching within the content of documents. It allows you to index the content of blobs stored in Azure Blob Storage, including PDFs, CSVs, Microsoft Office formats, and more. Cognitive Search extracts text and metadata from the documents, enabling advanced search functionality.

upvoted 1 times

 **neelkanths** 1 year, 2 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

answer is correct...

upvoted 2 times

 **raymond\_abcd** 1 year, 5 months ago

Answer is correct but keep in mind that Azure Cognitive Search is renamed to Azure AI Search. See: <https://learn.microsoft.com/en-us/azure/search/search-what-is-azure-search>

upvoted 4 times

 **Vukile** 1 year, 7 months ago

On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors

upvoted 7 times

 **dddddd111** 1 year, 7 months ago

Farmers and Distributors is that a new case study?

upvoted 6 times

 **Ciupaz** 1 year, 5 months ago

Yes, now is present in ExamTopics.

upvoted 2 times

🗨️ 👤 **RikinPatel** 1 year, 9 months ago

As per Goodle Bard this is correct ans.

Azure Blob index tags

Azure Cognitive Search

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 11 months ago

1. metadata

2. cognitive search

upvoted 2 times

🗨️ 👤 **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 3 times

🗨️ 👤 **Videira** 2 years, 4 months ago

On my exam 2023-02-25

upvoted 3 times

🗨️ 👤 **trini99** 2 years, 4 months ago

This question was on the 24/02/2023

upvoted 2 times

🗨️ 👤 **tmuras** 2 years, 4 months ago

This question was on exam on 17/02/23.

upvoted 3 times

🗨️ 👤 **Esward** 2 years, 5 months ago

Given answers are correct

Azure blob index tags

Cognitive serach

<https://learn.microsoft.com/en-us/azure/search/search-blob-storage-integration>

upvoted 4 times

🗨️ 👤 **ks1586** 2 years, 6 months ago

Did my exam on 17th December 2022. This question was on it.

upvoted 3 times

🗨️ 👤 **uditrm12** 2 years, 6 months ago

examtopic questions are enough to pass ?

upvoted 2 times

🗨️ 👤 **Ciupaz** 1 year, 9 months ago

Also follow some course on Udemy (or Pluralsight), just to see actions, and read all the discussions here on ExamTopics. Then, if you have time/energy, study Microsoft docs.

upvoted 2 times

🗨️ 👤 **HellBaby** 2 years, 6 months ago

No - you need a lot of other resources in order to get a better chance. The Microsoft exams are no joke.

upvoted 4 times

🗨️ 👤 **OPT\_001122** 2 years, 6 months ago

1. index tag

2. cognitive search

upvoted 2 times

## HOTSPOT -

You are developing a web application by using the Azure SDK. The web application accesses data in a zone-redundant BlockBlobStorage storage account.

The application must determine whether the data has changed since the application last read the data. Update operations must use the latest data changes when writing data to the storage account.

You need to implement the update operations.

Which values should you use? To answer, select the appropriate option in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

### Code evaluation

### Value

HTTP Header value

	▼
ETag	
Last Modified	
VersionId	

Conditional header

	▼
If-Match	
If-Modified-Since	
If-None-Match	

## Answer Area

### Code evaluation

### Value

HTTP Header value

	▼
ETag	
Last Modified	
VersionId	

Conditional header

	▼
If-Match	
If-Modified-Since	
If-None-Match	

Suggested Answer:

Box 1: Last Modified -

The Last-Modified response HTTP header contains a date and time when the origin server believes the resource was last modified. It is used as a validator to determine if the resource is the same as the previously stored one. Less accurate than an ETag header, it is a fallback mechanism.

Box 2: If-Modified-Since -

Conditional Header If-Modified-Since:

A DateTime value. Specify this header to perform the operation only if the resource has been modified since the specified time.

Incorrect:

Not ETag/If-Match -

Conditional Header If-Match:

An ETag value. Specify this header to perform the operation only if the resource's ETag matches the value specified. For versions 2011-08-18 and newer, the

ETag can be specified in quotes.

Reference:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Last-Modified> <https://docs.microsoft.com/en-us/rest/api/storageservices/specifying-conditional-headers-for-blob-service-operations>

  **tcybu**  2 years, 9 months ago

I think it should be:

- ETag - server returns this tag for a resource to ensure we operate on the same version of the resource in subsequent API calls
- If-Match - update is processed by the server only if the ETag provided matches the latest resource version ETag

The reason for that is we want to make sure we update the latest version of a resource:

"Update operations must use the latest data changes when writing"

So, when using Last-Modified with If-Modified-Since, the operation executes only when another client modifies the resource between our READ and WRITE operations.



If we wanted to use Last-Modified instead, we would need If-Unmodified-Since instead.

upvoted 45 times

  **Christian\_garcia\_martin** 11 months, 1 week ago

i agree , good explanation.

upvoted 1 times

  **Tenk14** 2 years, 9 months ago

I agree with you. (ETag + If-Match)

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/If-Match>

upvoted 4 times

  **[Removed]**  2 years, 9 months ago

I referred to the documentation also and I cannot convince my mind I care enough to even understand this because there is no chance I will ever have to know these options in details unless I have to use it. Who reads and understands all these information when this is available as documentation and with Microsoft documentation it is all about trial and error and hoping for the best

upvoted 24 times

  **ns4098**  9 months, 1 week ago

Based on the scenario where you need to implement update operations for a web application accessing data in a zone-redundant BlobStorage storage account, the application must ensure it uses the latest data by checking whether the data has changed since it was last read. Here are the appropriate choices:

1. HTTP Header value:

ETag

Reason: The ETag (Entity Tag) is a unique identifier assigned to a specific version of a resource. It is used to determine if the data has changed since the last request. When you read the data, you get the ETag, and when performing update operations, you can compare the ETag to ensure you are working with the latest version of the data.

2. Conditional header:

If-Match



Reason: The If-Match header is used to perform operations only if the ETag of the resource matches the one specified in the request. This ensures that updates happen only if the resource hasn't been modified since it was last read, preventing overwrites of more recent updates.

upvoted 2 times

  **richardc7** 1 year, 9 months ago

Got examn 28/09/23

upvoted 4 times

  **Nitin23** 1 year, 9 months ago

both are wrong in given solution.

Correct anser - etag and If-match

upvoted 2 times



🗄️ 👤 **LxstlnNsght** 1 year, 10 months ago

It's ETag and If-Match.

The ETag (Entity Tag) is an HTTP header that represents a unique identifier for a version of the resource. When sending an update request, you can include the If-Match header with the ETag value of the data you previously read. If the current ETag of the resource matches the provided ETag, it means the data hasn't changed.

upvoted 2 times

🗄️ 👤 **[Removed]** 1 year, 11 months ago

Correct answer

upvoted 1 times

🗄️ 👤 **calmriver44** 2 years, 1 month ago

It is Etag and If-Match

check this page <https://learn.microsoft.com/en-us/rest/api/storageservices/specifying-conditional-headers-for-blob-service-operations> for conditional header "If-Match" - We can specify ETag value. Specify this header to perform the operation only if the resource's ETag matches the value specified.

upvoted 2 times

🗄️ 👤 **adilkhan** 2 years, 3 months ago

Etag and if-match ChatGPT

upvoted 1 times

🗄️ 👤 **kandrew5** 2 years, 4 months ago

Since we also need the update to work which is a PUT operation the correct answer is ETag/If-Match

upvoted 1 times

🗄️ 👤 **mmdex** 2 years, 4 months ago

ETag

If-Match

There are 2 possible options to validate an update against unwanted changes:

- ETag + If-Match

- Last-Modified + If-Unmodified-Since

The second option is not available here, so it must be the first one. Besides, even if it was available, it would still be better to check ETag, because Last-Modified is weaker (it has only 1-second resolution, so it does not detect 2 changes within 1 second, so you cannot ensure you have the latest version) and I guess the zone-redundancy may also play some role in how Last-Modified is handled.

upvoted 3 times

🗄️ 👤 **tmndungu** 2 years, 4 months ago

For those who have taken AZ-204 test, should I trust ExamTopics answers or the Community's answers?

upvoted 5 times

🗄️ 👤 **rgullini** 2 years, 2 months ago

You should trust the discussions and other users experiences. The default answers in the questions might be right or wrong.

upvoted 7 times

🗄️ 👤 **Esward** 2 years, 5 months ago

Etag, If-Match are correct answers as per MS docs

Optimistic concurrency: An application performing an update will, as part of its update, determine whether the data has changed since the application last read that data. For example, if two users viewing a wiki page make an update to that page, then the wiki platform must ensure that the second update does not overwrite the first update. It must also ensure that both users understand whether their update was successful. This strategy is most often used in web applications.

<https://learn.microsoft.com/en-us/azure/storage/blobs/concurrency-manage?tabs=dotnet>

upvoted 4 times

🗄️ 👤 **Dotar** 2 years, 5 months ago

I think it should be E-tag and if-none-match.



See the following source on how e-tag works: [https://en.wikipedia.org/wiki/HTTP\\_ETag](https://en.wikipedia.org/wiki/HTTP_ETag)

upvoted 1 times

🗄️ 👤 **Skellecz** 2 years, 5 months ago

I think the answer is correct. The question is asking explicitly if "data changed since the application last read the data", so we must evaluate against the date, hence If-Modified-Since.

upvoted 2 times

  **AsuraX** 2 years, 6 months ago

Answer is correct - in my opinion.

- ETag - The ETag (or entity tag) HTTP response header is an identifier for a specific version of a resource. (Version!)

- Last-Modified - The Last-Modified response HTTP header contains a date and time when the origin server believes the resource was last modified. (Date and Time!)



(Source: <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers>)

In Question: The application must determine whether the data has changed since the application last read the data.

My understanding: We are working on dates, not versions! So, Last-Modified and If-Modified-Since are better options.

PS: I honestly have no idea why I need to know this for Azure Certification Exam :(

upvoted 7 times

  **capable** 2 years, 7 months ago

Got this question in the exam 27/11/2022, anyone, please mention the correct answer and explain

upvoted 1 times

## HOTSPOT -

An organization deploys a blob storage account. Users take multiple snapshots of the blob storage account over time.

You need to delete all snapshots of the blob storage account. You must not delete the blob storage account itself.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption  
snapshotsOption = Azure.Storage.Blobs.Models.

▼
DeleteIfExists
DeleteSnapshotsOption
WithSnapshot
WithSnapshotCore

▼
IncludeSnapshots
None
OnlySnapshots

## Suggested Answer:

## Answer Area

Delete (Azure.Storage.Blobs.Models.DeleteSnapshotsOption  
snapshotsOption = Azure.Storage.Blobs.Models.

▼
DeleteIfExists
DeleteSnapshotsOption
WithSnapshot
WithSnapshotCore

▼
IncludeSnapshots
None
OnlySnapshots

Box 1: DeleteSnapshotsOption -

Sample code in powershell:

//dont forget to add the include snapshots :)

```
await batchClient.DeleteBlobsAsync(listofURIforBlobs,
Azure.Storage.Blobs.Models.DeleteSnapshotsOption.IncludeSnapshots);
```

Sample code in .Net:


```
// Create a batch with three deletes
BlobBatchClient batchClient = service.GetBlobBatchClient();
BlobBatch batch = batchClient.CreateBatch();
batch.DeleteBlob(foo.Uri, DeleteSnapshotsOption.IncludeSnapshots); batch.DeleteBlob(bar.Uri, DeleteSnapshotsOption.OnlySnapshots);
batch.DeleteBlob(baz.Uri);
// Submit the batch
batchClient.SubmitBatch(batch);
```

Box 2: OnlySnapshots -

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/overview/azure/storage.blobs.batch-readme>

<https://stackoverflow.com/questions/39471212/programmatically-delete-azure-blob-storage-objects-in-bulks>

 **coffecold** Highly Voted 2 years, 2 months ago

it is DeleteSnapshotsOption.OnlySnapshots but the explanation in the answer is bogus as usual.

see <https://learn.microsoft.com/en-us/dotnet/api/azure.storage.blobs.models.deletesnapshotsoption?view=azure-dotnet>  
upvoted 20 times

 **micro9000** 2 years, 1 month ago

Yeah, we can ignore the explanation on this one,

We can use OnlySnapshots

upvoted 3 times

 **Esward** Highly Voted 1 year, 10 months ago

It was there in 13 Feb 2023 exam

upvoted 7 times

- 🗄️ 👤 **J\_obando13** Most Recent 1 month, 2 weeks ago  
Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)  
upvoted 1 times
- 🗄️ 👤 **harbox** 9 months, 2 weeks ago  
Is correct, <https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-delete>  
  
public static async Task DeleteBlobSnapshotsAsync(BlobClient blob)  
{  
// Delete a blob and all of its snapshots  
await blob.DeleteAsync(snapshotsOption: DeleteSnapshotsOption.IncludeSnapshots);  
  
// Delete only the blob's snapshots  
//await blob.DeleteAsync(snapshotsOption: DeleteSnapshotsOption.OnlySnapshots);  
}  
upvoted 3 times
- 🗄️ 👤 **Vukile** 1 year, 1 month ago  
On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors  
upvoted 2 times
- 🗄️ 👤 **chsiri** 11 months, 3 weeks ago  
Where can I find the case studies on this site.  
upvoted 1 times
- 🗄️ 👤 **Ciupaz** 11 months, 2 weeks ago  
Now is present.  
upvoted 1 times
- 🗄️ 👤 **Pulk** 1 year, 2 months ago  
correct.  
upvoted 1 times
- 🗄️ 👤 **MysticalSam** 1 year, 6 months ago  
This question was in today's exam on 10-June-2023  
upvoted 2 times
- 🗄️ 👤 **[Removed]** 1 year, 8 months ago  
Got this in exam today (5 April 2023)  
upvoted 4 times
- 🗄️ 👤 **Esward** 1 year, 11 months ago  
Given options are correct  
<https://learn.microsoft.com/en-us/dotnet/api/azure.storage.blobs.models.deletesnapshotoption?view=azure-dotnet>  
upvoted 1 times
- 🗄️ 👤 **OPT\_001122** 2 years ago  
DeleteSnapshotsOption.OnlySnapshots  
upvoted 4 times
- 🗄️ 👤 **finnishr** 2 years, 4 months ago  
Correct!!  
upvoted 5 times

## HOTSPOT -

An organization deploys a blob storage account. Users take multiple snapshots of the blob storage account over time.

You need to delete all snapshots of the blob storage account. You must not delete the blob storage account itself.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

## Answer Area

delete\_blob ( 

▼

delete\_container

delete\_snapshots

snapshot\_blob

snapshots\_present

 = 

▼

False

Include

Only

 )

Suggested Answer:

## Answer Area

delete\_blob ( 

▼

delete\_container

delete\_snapshots

snapshot\_blob

snapshots\_present

 = 

▼

False

Include

Only

 )

Box 1: delete\_snapshots -

# Delete only the snapshot (blob itself is retained)

blob\_client.delete\_blob(delete\_snapshots="only")

Box 2: only -



Reference:

[https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob\\_samples\\_common.py](https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py)

  **[Removed]** Highly Voted 2 years, 3 months ago

So, according to Microsoft a developer has to remember how the Microsoft guy chose to write in the .Net library and also how the python person thought and the unfortunate inconsistencies they built into these APIs?

upvoted 37 times

  **toysky731** 2 years ago

or you can retake the exam?

upvoted 4 times

  **bbq598** 2 years ago



haha I think you are right

upvoted 1 times

  **Jarynpl** 1 year, 1 month ago

I've read somewhere that before starting an exam you select in which language you are programming. Based on that you will get customize language questions/answers.

upvoted 3 times

  **ElectricPants** Highly Voted 2 years, 3 months ago

The answer looks correct and the link proves it:

Link to exact line with comment:

[https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob\\_samples\\_common.py#L65](https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py#L65)

upvoted 18 times

🗨️ 👤 **bgbgvfvf** Most Recent 1 year ago

I think given answer is correct.

upvoted 1 times

🗨️ 👤 **Pulk** 1 year, 2 months ago

given answers are correct.

upvoted 1 times

🗨️ 👤 **Esward** 1 year, 11 months ago

Given answers are correct

[https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob\\_samples\\_common.py](https://github.com/Azure/azure-sdk-for-python/blob/main/sdk/storage/azure-storage-blob/samples/blob_samples_common.py)

upvoted 1 times

🗨️ 👤 **gvasilev** 2 years ago

Ah yes my favorite thing as developer - to remember what is something doing instead of google it.. lmao nice "exam" micro\$oft

upvoted 5 times

🗨️ 👤 **Enigma\_\_\_** 2 years, 3 months ago

Correct

Required if the blob has associated snapshots. Values include:

"only": Deletes only the blobs snapshots.

"include": Deletes the blob along with all snapshots.

<https://learn.microsoft.com/en-us/azure/developer/python/sdk/storage/azure-storage-blob/azure.storage.blob.containerclient?view=storage-py-v12>

upvoted 4 times

🗨️ 👤 **LucasSomeone1** 2 years, 3 months ago

Weird comments. When there is only 1 or 2 comments, everyone is saying examtopics is correct...

upvoted 5 times

🗨️ 👤 **kampatra** 2 years, 3 months ago

Correct!

upvoted 2 times

## HOTSPOT

You are developing an application that monitors data added to an Azure Blob storage account.

You need to process each change made to the storage account.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

```
var changeFeedClient = new BlobServiceClient("...").GetChangeFeedClient();
var x = default(string);
while (true)
{
    var changeFeed = changeFeedClient.
    foreach (var c in changeFeed)
    {
        x = c.
        ProcessChanges(c.Values);
    }
}
```

Dropdown 1 options: GetChanges(), GetChangesAsync(), GetChanges(x).AsPages(), GetChanges(x).GetEnumerator()

Dropdown 2 options: ContinuationToken, GetRawResponse().ReasonPhrase, Values.Max(x => x.EventTime).ToString(), Values.Min(x => x.EventTime).ToString()

## Answer Area

```
var changeFeedClient = new BlobServiceClient("...").GetChangeFeedClient();
var x = default(string);
while (true)
{
    var changeFeed = changeFeedClient.
    foreach (var c in changeFeed)
    {
        x = c.
        ProcessChanges(c.Values);
    }
}
```

Dropdown 1 options: GetChanges(), GetChangesAsync(), GetChanges(x).AsPages(), GetChanges(x).GetEnumerator()

Dropdown 2 options: ContinuationToken, GetRawResponse().ReasonPhrase, Values.Max(x => x.EventTime).ToString(), Values.Min(x => x.EventTime).ToString())

## Suggested Answer:

 tcybu  2 years, 5 months ago

I've tried the code in VS. Here's some thoughts:

1. box:

- GetChanges() - wrong - var c in the foreach would be BlobChangeFeedEvent which doesn't contain Values property used in ProcessChanges(c.Values) line below
- GetChangesAsync() - wrong - code won't compile because it would require await foreach loop instead
- GetChanges(x).AsPages() - correct - it's the only option to make this code even compile
- GetChanges(x).GetEnumerator() - wrong - you cannot use IEnumerable type as foreach source

2. box:



- `x = c.ContinuationToken` - right - variable `x` was used as `continuationToken` parameter in `changeFeedClient.GetChanges(x).AsPages()` above
- `c.GetRawResponse().ReasonPhrase` - wrong - that does not make sense to use this value as continuation token
- `x = c.Values.Min` - wrong - continuation token is a number not date
- `x = c.Values.Max` - wrong - as above

So to sum up

1. `changeFeedClient.GetChanges(x).AsPages()`
2. `x = c.ContinuationToken;`

You can find more about Continuation Token here:



<https://jessehouwing.net/azure-devops-accessing-apis-with-large-volumes-of-data/>  
upvoted 83 times

  **applepie** 1 year, 11 months ago

got this question today, go with this answer - 7/30/2023, score 895/1000  
upvoted 3 times

  **Elbanna** 1 year, 9 months ago

Does all questions came from exam topics ?  
upvoted 2 times

  **Net\_IT** 1 year, 8 months ago

The majority of the questions  
upvoted 1 times

  **DP\_Bhatt** 2 years, 5 months ago

Agree.

Box 1. `GetChanges(x).AsPages()`

Box 2. `ContinuationToken`

upvoted 5 times

  **Michael2023** 2 years, 5 months ago

you are right...

1. `changeFeedClient.GetChanges(x).AsPages()` -> returns an `IEnumerable<Page<BlobChangeFeedEvent>>` ... when you loop through these pages "Page<BlobChangeFeedEvent>" you will get the options "page.ContinuationToken" and `page.Values` which are used in this example  
100% Correct...

found the code here

<https://github.com/Azure/azure-sdk-for-net/tree/main/sdk/storage/Azure.Storage.Blobs.ChangeFeed#resume-with-continuationtoken>

upvoted 4 times

  **warchoon** 2 years, 4 months ago

`var x = default(string);`

so, `x` is string

`ContinuationToken` is not string

upvoted 3 times

  **vizay**  2 months ago


`BlobServiceClient` is the main tool you use to interact with your blob storage.

You use `GetChangeFeedClient` to monitor changes in your blobs.

`GetChangesAsPage` allows you to fetch changes in smaller chunks (pages), especially when there are too many.

The `ContinuationToken` helps you continue from where you left off in case you can't process all changes in one go.

upvoted 1 times

  **4bd3116** 10 months, 1 week ago

The Answer is correct!!

```
var changeFeedClient = new BlobServiceClient("..").GetChangeFeedClient();
```

```
var continuationToken = default(string);
```



```
while (true)
{
    var changeFeed = changeFeedClient.GetChanges(continuationToken);
    foreach (var change in changeFeed)
    {
        continuationToken = change.ContinuationToken;
        ProcessChanges(change.Values);
    }
}
```

<https://learn.microsoft.com/en-us/dotnet/api/azure.storage.blobs.changefeed.blobchangefeedclient.getchanges?view=azure-dotnet-preview#azure-storage-blobs-changefeed-blobchangefeedclient-getchanges>

upvoted 1 times

🗳️ 👤 **8ac3742** 10 months, 4 weeks ago

The answer is incorrect, continuationToken points to the position of the last processed change event, changeFeedClient.Getchanges(x).AsPages retrieves changes in chunks or pages so it's more efficient Handling.

upvoted 1 times

🗳️ 👤 **Stel0Papad4** 1 year, 1 month ago

For those who are wondering of the answer see this link by following the methods you see its using above above.

1. GetChanges(x).AsPages()
2. ContinuationToken

is the correct answer

upvoted 1 times

🗳️ 👤 **Stel0Papad4** 1 year, 1 month ago

<https://learn.microsoft.com/en-us/dotnet/api/azure.storage.blobs.changefeed.blobchangefeedclient.getchanges?view=azure-dotnet-preview#azure-storage-blobs-changefeed-blobchangefeedclient-getchanges>

upvoted 1 times

🗳️ 👤 **neelkanths** 1 year, 2 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

- GetChanges(x).AsPages()

ContinuationToken;

upvoted 3 times

🗳️ 👤 **Nihilist11** 1 year, 5 months ago

Jan-21-2024 - This was in my exam - Score 740, just border :) Didnt purchased contributor access.

Contosso case study [Couldnt find here]

upvoted 2 times

🗳️ 👤 **1989jay** 1 year, 5 months ago

This question was on exam. 09/01/2024. Passed 872.

upvoted 2 times

🗳️ 👤 **leviata5** 1 year, 9 months ago

Got this quesiton in examn - 2023.09.25. Got Case Study Contoso

upvoted 2 times

🗳️ 👤 **kajal2205** 1 year, 6 months ago

What answer you selected?

upvoted 1 times

🗳️ 👤 **NightshadeRC** 1 year, 11 months ago

Had this question in today's exam: 2023-07-26

upvoted 1 times

🗳️ 👤 **Videira** 2 years, 4 months ago

On my exam 2023-02-25

upvoted 2 times

🗳️ 👤 **tmuras** 2 years, 4 months ago

This question was on exam on 17/02/23.

upvoted 1 times

🗨️ 👤 **red0400** 2 years, 4 months ago

what answer do you put ?

upvoted 1 times

🗨️ 👤 **Esward** 2 years, 5 months ago

As per Microsoft docs:

GetchangesAsync(), continuationtoken are correct

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed-how-to>

upvoted 2 times

🗨️ 👤 **drjakiell** 2 years, 5 months ago

for GetchangesAsync you need the async await approach

upvoted 3 times

🗨️ 👤 **AfPham** 2 years, 4 months ago

Async option can be ignore, it requires 'await', which is not there in this situation

upvoted 5 times

🗨️ 👤 **frks** 2 years, 5 months ago

I believe it should be await... getchangesasync.... then getrawresponse, as x is a string. I don't see the connection to cancellationtoken.

upvoted 1 times

🗨️ 👤 **chettir01** 2 years, 5 months ago

i think the first box must be : getChangesAsync()

the second one is correct

source : <https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-change-feed-how-to>

upvoted 3 times

🗨️ 👤 **MikeSantos** 2 years, 2 months ago

I don't agree. There is no await statement before the method

upvoted 2 times

## HOTSPOT

-

You develop an application that sells AI generated images based on user input. You recently started a marketing campaign that displays unique ads every second day.

Sales data is stored in Azure Cosmos DB with the date of each sale being stored in a property named 'whenFinished'.

The marketing department requires a view that shows the number of sales for each unique ad.

You need to implement the query for the view.

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

NOTE: Each correct selection is worth one point.

SELECT

max(c.whenFinished)  
sum(c.whenFinished)  
count(c.whenFinished)

FROM c

group by

DateTimeBin(c.whenFinished, 'day', 2)  
DateTimePart(c.whenFinished, 'day', 2)  
DateTimeBin(c.whenFinished, 'hour', 12)  
DateTimePart(c.whenFinished, 'hour', 12)

Suggested Answer:

SELECT

max(c.whenFinished)  
sum(c.whenFinished)  
count(c.whenFinished)

DateTimeBin(c.whenFinished, 'day', 2)  
DateTimePart(c.whenFinished, 'day', 2)  
DateTimeBin(c.whenFinished, 'hour', 12)  
DateTimePart(c.whenFinished, 'hour', 12)

FROM c

group by

DateTimeBin(c.whenFinished, 'day', 2)  
DateTimePart(c.whenFinished, 'day', 2)  
DateTimeBin(c.whenFinished, 'hour', 12)  
DateTimePart(c.whenFinished, 'hour', 12)

Correct!




Can't be DateTimePart as it takes two args only, see <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/date-time-functions>  
upvoted 26 times

  **Dotar** 2 years, 5 months ago

You are right good sir I stand corrected.  
upvoted 2 times

  **MysticalSam**  2 years ago

This question was in today's exam on 10-June-2023  
upvoted 6 times

  **J\_obando13**  1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)  
upvoted 2 times

  **f09229d** 3 months ago

In exam 30/3/2025 with this answer, I scored 907/1000  
upvoted 1 times

  **38195c0** 1 year ago

"The marketing department requires a view that shows the number of sales for each unique ad." - where is the ad in the query?  
upvoted 2 times

  **DamuKeesh** 5 months ago



Cosider every 2nd day is generating unique ad  
upvoted 1 times

  **overhill** 8 months ago


exactly, wrong question or wrong image  
upvoted 1 times

  **sekelsenmat** 1 year, 4 months ago

On exam 2024, went with given answer, score 872 or something. Case Study: Farmers and Distributors  
upvoted 4 times

  **Vukile** 1 year, 7 months ago



On exam 9 Nov 2023, went with given answer, socre 865. Case Study: Farmers and Distributors  
upvoted 4 times

  **Ciupaz** 1 year, 7 months ago



This question was inside that Case Study?  
upvoted 1 times

  **richardc7** 1 year, 9 months ago

got today in 28-09-23  
upvoted 1 times

  **janzak** 2 years, 2 months ago

why count and not sum?  
upvoted 2 times

  **kggy01** 2 years, 1 month ago

[...]The marketing department requires a view that shows the "number of sales"[...]  
upvoted 1 times

  **MDWPartners** 1 year, 9 months ago

the field is a date, not a number, so you cannot sum it.  
upvoted 1 times

  **Esward** 2 years, 5 months ago

Given answers Count, DateTimeBin are correct answers as per MS docs  
DateTimebin takes 3 and 4 parameters but DateTimepart takes only 2 args

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/date-time-functions>  
upvoted 4 times

🗨️ 👤 **DP\_Bhatt** 2 years, 5 months ago

Correct!

upvoted 1 times

🗨️ 👤 **Dotar** 2 years, 5 months ago

Anyone have an idea as to why it says day and not hour?

upvoted 2 times

🗨️ 👤 **plusJoyed** 2 years, 1 month ago

It says every second day, so 12 hours would be twice a day. It might work if hours was 48

upvoted 1 times

🗨️ 👤 **Dotar** 2 years, 5 months ago

Pretty sure it has to be DateTimePart(...day) option.

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/query/datetimetypepart>

upvoted 1 times

## HOTSPOT

-

You implement an Azure solution to include Azure Cosmos DB, the latest Azure Cosmos DB SDK, and the Core (SQL) API. You also implement a change feed processor on a new container instance by using the Azure Functions trigger for Azure Cosmos DB.

A large batch of documents continues to fail when reading one of the documents in the batch. The same batch of documents is continuously retried by the triggered function and a new batch of documents must be read.

You need to implement the change feed processor to read the documents.

Which feature should you implement? To answer, select the appropriate features in the answer area.

NOTE: Each correct selection is worth one point.

**Requirement**

Read a new batch of documents while keeping track of the failing batch of documents.

Handle errors in the change feed processor.

**Feature**

Lease container  
Dead-letter queue  
Life-cycle notifications  
Change feed estimator

Lease container  
Dead-letter queue  
Life-cycle notifications  
Change feed estimator

**Requirement**

Read a new batch of documents while keeping track of the failing batch of documents.



**Suggested Answer:**

Handle errors in the change feed processor.

**Feature**

Lease container  
Dead-letter queue  
Life-cycle notifications  
**Change feed estimator**

Lease container  
**Dead-letter queue**  
Life-cycle notifications  
Change feed estimator

 **sarmaria**  1 year, 9 months ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 14 times

 **Mattfusa** 10 months, 1 week ago

For case study's though dont they block by contributor access now? How did you prepare for those?

upvoted 1 times

 **ReyPirata** 1 year, 4 months ago


Agreed.

This was on the exam (08/20/2023). Scored 925

Change feed estimator

Dead-letter queue

upvoted 5 times

 **imanonion**  1 year, 11 months ago

I agree with the answer!

1. Change feed estimators monitor change feed processors' progress



2. Dead-letter queues handle errors and are able to monitor failed attempts, require failed attempts and even trigger a follow-up action (remediation or response)

upvoted 11 times

  **fuji36850** Most Recent 7 months, 1 week ago



Why feed estimator ? It has nothing to do with 'keeping track of failing batch'. Failing batch should be sent to dead-letter queue to keep track of it...

upvoted 1 times

  **azpardha** 8 months, 3 weeks ago

Got this on 04/07/2024, Went with highly voted answer, score:820

upvoted 2 times

  **Hendrikdb** 10 months, 2 weeks ago

The correct answer seems to me:

- Change feed estimator

- Life cycle notifications

You also can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed, or you can use life cycle notifications to detect underlying failures.

Life cycle notifications literally says:

Register a handler for `WithErrorNotification` to be notified when the current host encounters an exception during processing. You need to be able to distinguish whether the source is the user delegate (an unhandled exception) or an error that the processor encounters when it tries to access the monitored container (for example, networking issues).



upvoted 3 times

  **Hendrikdb** 10 months, 2 weeks ago

Correction, change feed estimator does nothing in this question. I do not know why people pick this answer and I blindly followed. Change feed estimator only tracks changes between read operations.

Life cycle notifications tracks everything. and you can use a dead letter queue to process the failed documents.

upvoted 3 times

  **Woksi** 10 months, 3 weeks ago

Got on 07/02/2024

upvoted 2 times

  **Ciupaz** 1 year, 1 month ago

The answers should be reversed.

upvoted 3 times

  **AndySmith** 1 year, 1 month ago

On exam 3-Nov-2023. Went with proposed answer - 932/1000.

1) change feed estimator

2) dead-letter queue

upvoted 2 times

  **dom271219** 1 year, 3 months ago

1/ "Change feed estimator" solely estimates the rate of changes, number of shards, and distribution of the feed, but not read or process documents. There's a error because by confusing with "Change feed processor".

2/ Dead-letter queue

upvoted 2 times

  **[Removed]** 1 year, 5 months ago

should be vice versa

upvoted 4 times

  **juanckar** 1 year, 5 months ago

This was on the exam (July 2023). Unsure about solution....

upvoted 2 times

  **aragones** 1 year, 7 months ago

Got this on 2023may12.

my cases also:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 2 times

🗨️ 👤 **Defa** 1 year, 9 months ago

From: <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#error-handling>

"To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue. This design ensures that you can keep track of unprocessed changes while still being able to continue to process future changes. The errored-message queue might be another Azure Cosmos DB container. The exact data store does not matter, simply that the unprocessed changes are persisted.

In addition, you can use the change feed estimator to monitor the progress of your change feed processor instances as they read the change feed or use the life cycle notifications to detect underlying failures."

As the questions says 'keeping track' I'd go for feed estimator.

upvoted 3 times

🗨️ 👤 **Videira** 1 year, 10 months ago

On my exam 2023-02-25

upvoted 2 times

🗨️ 👤 **Ayman99** 1 year, 10 months ago

Was on the exam 27.02.2023

Went for Lifecycle notifications and dead letter queue.

Score 870

upvoted 3 times

🗨️ 👤 **mabdo** 1 year, 11 months ago

was on exam 02/2023

upvoted 3 times

🗨️ 👤 **alekomar** 1 year, 11 months ago

1. Life-cycle notifications

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#life-cycle-notifications>

Register a handler for WithErrorNotification to be notified when the current host encounters an exception during processing

2. Dead-letter queue

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#error-handling>

To prevent your change feed processor from getting "stuck" continuously retrying the same batch of changes, you should add logic in your delegate code to write documents, upon exception, to an errored-message queue.

upvoted 4 times

🗨️ 👤 **warchoon** 1 year, 10 months ago

Correct

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-processor?tabs=dotnet#error-handling>:~:text=use%20the%20life%20cycle%20notifications%20to%20detect%20underlying%20failures.

upvoted 2 times

🗨️ 👤 **warchoon** 1 year, 10 months ago

but

1. (about the stuck) - Dead-letter and 2. (about error handling) - notifications

upvoted 4 times



You are developing an application to store business-critical data in Azure Blob storage.

The application must meet the following requirements:

- Data must not be modified or deleted for a user-specified interval.
- Data must be protected from overwrites and deletes.
- Data must be written once and allowed to be read many times.

You need to protect the data in the Azure Blob storage account.

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

NOTE: Each correct selection is worth one point.

- A. Configure a time-based retention policy for the storage account.
- B. Create an account shared-access signature (SAS).
- C. Enable the blob change feed for the storage account.
- D. Enable version-level immutability support for the storage account.
- E. Enable point-in-time restore for containers in the storage account.
- F. Create a service shared-access signature (SAS).

**Suggested Answer: AF**

Community vote distribution

AD (85%)

Other

 **linhbg** Highly Voted 2 years, 5 months ago

I think the answer should be AD

A. Configure a time-based retention policy for the storage account

- A time-based retention policy stores blob data in a Write-Once, Read-Many (WORM) format for a specified interval. When a time-based retention policy is set, clients can create and read blobs, but can't modify or delete them. After the retention interval has expired, blobs can be deleted but not overwritten.

<https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-time-based-retention-policy-overview>

D. Before you can apply a time-based retention policy to a blob version, you must enable support for version-level immutability.

<https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal>

upvoted 39 times

 **Esward** 2 years, 5 months ago


I Agree with you A and D are correct answers

upvoted 3 times

 **surprise0011** 2 years, 2 months ago

received 2023-04-17 went with above answer, score 926

upvoted 5 times

 **Vichu\_1607** Most Recent 7 months, 3 weeks ago

**Selected Answer: AD**

A. Configure a time-based retention policy for the storage account.

D. Enable version-level immutability support for the storage account.

upvoted 1 times

 **murat12345** 8 months, 4 weeks ago

a: Time-based retention policies: With a time-based retention policy, users can set policies to store data for a specified interval. When a time-based retention policy is set, objects can be created and read, but not modified or deleted. After the retention period has expired, objects can be deleted but not overwritten.

d: Specify whether to enable version-level immutability for this container. Version-level immutability can be applied to specific blobs (any or all) in the container. For blobs with version-level immutability set, blob overwrites will still be allowed, but Azure will maintain immutable versions of each blob. Once enabled, this setting cannot be removed. Versioning is required for this feature, and cannot be disabled on a storage account while version-level policies are in place

Not sure why D is needed, because with A you already cannot add or change the data.

upvoted 1 times

🗨️ 👤 **8ac3742** 10 months, 4 weeks ago

The answer is A and D: retention can prevent modification and deletion during a period, and version level immutability can prevent each version from being overwritten and deleted. For some frequently changed data, the user cannot apply a retention policy, the user can only set version-level immutability.

upvoted 1 times

🗨️ 👤 **Pulk** 1 year, 8 months ago

correct answer is AD

upvoted 1 times

🗨️ 👤 **AashirwadSinha** 1 year, 10 months ago

A and D are the correct answers

upvoted 1 times

🗨️ 👤 **andrewnm** 1 year, 11 months ago

**Selected Answer: AD**

A is clearly a correct answer and this documentation confirms <https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-policy-configure-version-scope?tabs=azure-portal> that D is also required.

upvoted 1 times

🗨️ 👤 **deathRac3** 2 years, 1 month ago

**Selected Answer: AF**

Feels correct

upvoted 1 times

🗨️ 👤 **aragones** 2 years, 1 month ago

Got this type 2023-05-12.

My case:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data

case study: VanArsdel Inc Canada

upvoted 2 times

🗨️ 👤 **adilkhan** 2 years, 3 months ago

A,D chatgpt

upvoted 2 times

🗨️ 👤 **Saluk\_DE** 2 years, 3 months ago

Question was in Exam 2023-03-30

upvoted 1 times

🗨️ 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Went with AD Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

🗨️ 👤 **Ayman99** 2 years, 4 months ago

Was on the exam 27.02.2023

Went for A and D.

Score 870

upvoted 2 times

🗨️ 👤 **trini99** 2 years, 4 months ago

Which resources did you use to study?

upvoted 2 times

🗨️ 👤 **uffuchsi** 2 years, 4 months ago

**Selected Answer: AD**

100% A and D, other options don't fit into this question (i.e., don't solve it)

upvoted 1 times

🗨️ 👤 **mabdo** 2 years, 4 months ago

on exam 02/2023

upvoted 1 times

🗨️ 👤 **alexsein74** 2 years, 5 months ago

**Selected Answer: AD**

A. Configuring a time-based retention policy for the storage account would ensure that the data in the storage account cannot be modified or deleted for a user-specified interval. This would meet the requirement that data must not be modified or deleted for a user-specified interval.

D. Enabling version-level immutability support for the storage account would ensure that the data in the storage account is protected from overwrites and deletes. This would meet the requirement that data must be protected from overwrites and deletes, and written once and allowed to be read many times.

upvoted 1 times

🗨️ 👤 **DP\_Bhatt** 2 years, 5 months ago

**Selected Answer: AD**

Should be A and D

upvoted 2 times

You are updating an application that stores data on Azure and uses Azure Cosmos DB for storage. The application stores data in multiple documents associated with a single username.

The application requires the ability to update multiple documents for a username in a single ACID operation.

You need to configure Azure Cosmos DB.

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

NOTE: Each correct selection is worth one point.

- A. Create a collection sharded on username to store documents.
- B. Configure Azure Cosmos DB to use the Gremlin API.
- C. Create an unsharded collection to store documents.
- D. Configure Azure Cosmos DB to use the MongoDB API.

**Suggested Answer:** CD

Community vote distribution

CD (100%)

🗳️ 👤 **imanonion** Highly Voted 1 year, 11 months ago

**Selected Answer:** CD

Agreed!

To update multiple documents for a username in a single ACID operation in Azure Cosmos DB, you need to ensure that the documents are stored in the same logical partition.

To do this, you should perform the following actions:

Create an unsharded collection to store documents. This will ensure that all documents are stored in the same logical partition. Configure Azure Cosmos DB to use the MongoDB API. The MongoDB API supports multi-document ACID transactions, which allow you to update multiple documents in a single atomic operation.

upvoted 26 times

🗳️ 👤 **Dotar** 1 year, 11 months ago

How do you know that the documents should be stored in the same logical partition. Is this a requirement for MongoDB API?

upvoted 3 times

🗳️ 👤 **surprise0011** 1 year, 8 months ago

all operations like executing stored procedures or udf are effecting only items within single partition regardless of used API

upvoted 1 times

🗳️ 👤 **[Removed]** Highly Voted 1 year, 8 months ago

Got this in exam today (5 April 2023)

upvoted 6 times

🗳️ 👤 **laminall** Most Recent 1 month, 4 weeks ago

**Selected Answer:** AD

According to Elon AI

upvoted 1 times

🗳️ 👤 **raymond\_abcd** 11 months, 1 week ago

Forget previous comment, not correct

upvoted 2 times

🗳️ 👤 **raymond\_abcd** 11 months, 1 week ago

I think it is AD because you need to shard across partitions on username to be able to process all documents with that user name in one transactional batch. You have to create a shard key for this in the cosmos db container based on the username.

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 5 months ago

Correct answer is AD

upvoted 3 times

🗨️ 👤 **Bokomi** 1 year, 7 months ago

**Selected Answer: CD**

To update multiple documents for a username in a single ACID operation in Azure Cosmos DB, you should perform the following actions:

C. Create an unsharded collection to store documents, since sharded collections do not support multi-document transactions.

D. Configure Azure Cosmos DB to use the MongoDB API, since it is the only API that supports multi-document transactions.

Therefore, the correct actions are C and D.

upvoted 4 times

🗨️ 👤 **Xin0613** 1 year, 8 months ago

Agreed!

upvoted 1 times

🗨️ 👤 **Esward** 1 year, 10 months ago

It was there in 13 Feb 2023 exam

upvoted 3 times

🗨️ 👤 **Esward** 1 year, 11 months ago

CD are correct

update multiple documents for a username in a single ACID operation.

upvoted 3 times

🗨️ 👤 **Esward** 1 year, 11 months ago

<https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/use-multi-document-transactions#requirements>

upvoted 5 times

🗨️ 👤 **g2000** 1 year, 11 months ago

**Selected Answer: CD**

multi-document transaction is only supported within unsharded collection.

<https://learn.microsoft.com/en-us/azure/cosmos-db/mongodb/use-multi-document-transactions#requirements>

upvoted 3 times

You develop Azure solutions.

You must connect to a No-SQL globally-distributed database by using the .NET API.

You need to create an object to configure and execute requests in the database.

Which code segment should you use?

- A. `database_name = 'MyDatabase'`  
`database = client.create_database_if_not_exists(id=database_name)`
- B. `client = CosmosClient(endpoint, key)`
- C. `container_name = 'MyContainer'`  
`container = database.create_container_if_not_exists(`  
`id=container_name, partition_key=PartitionKey(path="/lastName"), offer_throughput=400 )`

**Suggested Answer: C**

*Community vote distribution*

B (100%)

 **imanonion** Highly Voted 2 years, 5 months ago

**Selected Answer: B**

CosmosClient has to be created before you can do option A and C to create databases and execute requests.

```
client = CosmosClient(endpoint, key)
database_name = 'MyDatabase'
database = client.create_database_if_not_exists(id=database_name)
container_name = 'MyContainer'
container = database.create_container_if_not_exists(
id=container_name, partition_key=PartitionKey(path="/lastName"), offer_throughput=400 )
```

upvoted 19 times

 **[Removed]** 1 year, 11 months ago

It doesn't matter. Question is asking about .NET API. A and C are not .NET code

upvoted 9 times

 **examtopicsLogin123** Most Recent 9 months, 2 weeks ago

I found that this question has the following answer options:

- A. `new Container(EndpointUri, PrimaryKey);`
- B. `new Database(EndpointUri, PrimaryKey);`
- C. `new CosmosClient(EndpointUri, PrimaryKey);`

and the correct answer is: `new CosmosClient(...)`

upvoted 1 times

 **tempacc4nk** 1 year, 10 months ago

I think its B, but its missing new keyword.

```
CosmosClient cosmosClient = new CosmosClient(
"connection-string-from-portal",
new CosmosClientOptions()
{
ApplicationRegion = Regions.EastUS2,
});
```

upvoted 4 times

🗳️ 👤 **LHU** 8 months ago

Agreed, all these examples are incorrect C# code. Also note the strings denoted as '' and not ""  
upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 11 months ago

B is only one .NET code here. So, B is a correct answer  
upvoted 3 times

🗳️ 👤 **kabbas** 2 years ago

**Selected Answer: B**

You need to create the client before anything else  
upvoted 4 times

🗳️ 👤 **deathRac3** 2 years, 1 month ago

**Selected Answer: B**

correct  
upvoted 1 times

🗳️ 👤 **rcardenas72** 2 years, 4 months ago

The correct code segment to configure and execute requests in the No-SQL globally-distributed database by using the .NET API would be option B:  
client = CosmosClient(endpoint, key)  
This code initializes the CosmosClient class and connects to the database with the specified endpoint and key parameters. This is the first step in creating a client to communicate with the Azure Cosmos DB.

Option A shows how to create a database, and option C shows how to create a container inside a database, but they do not create the client object that is needed to communicate with the database. Therefore, they are not correct options for this scenario.  
upvoted 4 times

🗳️ 👤 **uffuchsi** 2 years, 4 months ago

Has to be B - How could you do either A or C before B?  
upvoted 2 times

🗳️ 👤 **BrettusMaximus** 2 years, 4 months ago

A  
You need to "create an object" to configure and execute requests in the "database".  
Not B  
Sure B must be done before A, but A is the actual database Object  
Not C  
C is the container and not the database object  
upvoted 2 times

🗳️ 👤 **LeoIT** 2 years, 5 months ago

B as "by using the .NET API"  
upvoted 2 times

🗳️ 👤 **Esward** 2 years, 5 months ago

B is the correct answer  
<https://learn.microsoft.com/en-us/python/api/azure-cosmos/azure.cosmos.cosmosclient?view=azure-python>  
upvoted 1 times

🗳️ 👤 **AmitPaul** 2 years, 5 months ago

**Selected Answer: B**

CosmosClient has to be created first  
upvoted 2 times

🗳️ 👤 **Kavin235** 2 years, 5 months ago

**Selected Answer: B**



I believe Cosmos client type object can be used to configure and execute requests in the Azure Cosmos DB database service.  
<https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.cosmosclient?view=azure-dotnet>  
upvoted 1 times

🗳️ 👤 **Skelletcz** 2 years, 5 months ago

**Selected Answer: B**

I vote for CosmosClient, which can connect to No-SQL database

upvoted 3 times

  **Nhiendo** 2 years, 5 months ago

i think answer B

upvoted 2 times

  **chettir01** 2 years, 5 months ago

i don't think so, in my point of view, cosmos client can be used to create your datababe, but not to manipulate it.  
to do that, you have to create a container from this database, and use this object to do some requests !

upvoted 1 times



You develop a web application that provides access to legal documents that are stored on Azure Blob Storage with version-level immutability policies. Documents are protected with both time-based policies and legal hold policies. All time-based retention policies have the AllowProtectedAppendWrites property enabled.

You have a requirement to prevent the user from attempting to perform operations that would fail only when a legal hold is in effect and when all other policies are expired.

You need to meet the requirement.

Which two operations should you prevent? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. adding data to documents
- B. deleting documents
- C. creating documents
- D. overwriting existing documents

**Suggested Answer: BD**

Community vote distribution

BD (100%)

 **minYoongi** Highly Voted 2 years, 2 months ago

The legal hold policies prevent the deletion of the legal documents that are stored on Azure Blob Storage, but they do not prevent other types of operations such as read, write, or update. Therefore, you need to prevent the following two operations that would fail only when a legal hold is in effect and when all other policies are expired:


Delete operation: When a legal hold is in effect and all other policies are expired, attempting to delete the document will fail. Therefore, you should prevent the user from attempting to delete a document.

Overwrite operation: When a legal hold is in effect and all other policies are expired, attempting to overwrite the document with a new version will fail. Therefore, you should prevent the user from attempting to overwrite a document with a new version.

By preventing these two operations, you can ensure that the legal documents are not accidentally deleted or overwritten when they are protected by legal hold policies.

CHatgpt answer. Is this correct?

upvoted 14 times

 **mihailos** 1 year, 10 months ago


Got it in exam 28/08/23. Went with this answer. Scored 912

upvoted 3 times

 **Ivanlco** 2 years ago

ty for the comment made it much clearer

upvoted 1 times

 **juanckar** Highly Voted 1 year, 11 months ago

This was on the exam (July 2023). Went with BD. Scored 917

upvoted 7 times

 **b1992** 1 year, 11 months ago

What was your casestudy?

upvoted 1 times

 **macobuzi** 1 year, 10 months ago

What does it mean casestudy?

upvoted 1 times

🗲️ 👤 **ce77bdc** Most Recent 1 year, 1 month ago

What's confusing is 'adding data to documents' sure sounds like Updating documents...

upvoted 4 times

🗲️ 👤 **jobolesonihal** 1 year, 3 months ago

Selected Answer: BD

When a legal hold is applied to a container, all existing blobs move into an immutable WORM state in less than 30 seconds. All new blobs that are uploaded to that policy-protected container will also move into an immutable state. Once all blobs are in an immutable state, overwrite or delete operations in the immutable container aren't allowed.

upvoted 1 times

🗲️ 👤 **Woksi** 1 year, 4 months ago

Got on 07/02/2024

upvoted 1 times

🗲️ 👤 **AndySmith** 1 year, 7 months ago

On exam 3-Nov-2023. Wend with proposed answer - 932/1000.

B & D

upvoted 4 times

🗲️ 👤 **Ciupaz** 1 year, 6 months ago

How do you remember all questions?

upvoted 1 times

🗲️ 👤 **training\_bit2bit** 1 year, 9 months ago

legal hold policies: prevents delete and overwrite actions.

the question said that "legal hold" is in effect an the others expired, we should prevent other operations from user. so the answer should not be: A and C ?

B and D are already prevented with the legal hold policies, right?

upvoted 1 times

🗲️ 👤 **[Removed]** 1 year, 11 months ago

I don't really understand this question

upvoted 5 times

🗲️ 👤 **Christian\_garcia\_martin** 10 months, 3 weeks ago

is the way is written , it happen with most of them

upvoted 1 times

🗲️ 👤 **adilkhan** 2 years, 2 months ago

b,d IS CORRECT

upvoted 3 times

🗲️ 👤 **halfway** 2 years, 2 months ago

Selected Answer: BD

BD. Ref: <https://learn.microsoft.com/en-us/azure/storage/blobs/immutable-legal-hold-overview>

upvoted 5 times

HOTSPOT

-

You provisioned an Azure Cosmos DB for NoSQL account named account1 with the default consistency level.

You plan to configure the consistency level on a per request basis. The level needs to be set for consistent prefix for read and write operations to account1.

You need to identify the resulting consistency level for read and write operations.

Which levels should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

#### Operation type

Read operations

#### Resulting consistency level

▼

strong  
session  
consistent prefix

Write operations

▼

strong  
session  
consistent prefix

### Answer Area

Suggested Answer:

#### Operation type

Read operations

#### Resulting consistency level

▼

strong  
session  
**consistent prefix**

Write operations

▼

strong  
**session**  
consistent prefix

 **adilkhan**  2 years, 2 months ago

Just passed the exam 26/04/2023. All the questions were from Exam topics. Got 970 marks!

upvoted 31 times

 **[Removed]** 1 year, 11 months ago

useless comment

upvoted 19 times

🗨️ 👤 **LucreziaB** 1 year, 11 months ago

Still more useful than yours :)

upvoted 42 times

🗨️ 👤 **Christian\_garcia\_martin** 11 months, 1 week ago

congrats i got 984 in AZ-900 but i don't think i will get similar score in AZ-204

upvoted 1 times

🗨️ 👤 **UzrShkl** 2 years, 1 month ago

Bro you had the contributor access?

upvoted 5 times

🗨️ 👤 **Christian\_garcia\_martin** 11 months, 1 week ago

you can't get 970 only with the 100 first questions.

upvoted 1 times

🗨️ 👤 **Rodikito** 1 year, 10 months ago

of course

upvoted 1 times

🗨️ 👤 **Ciupaz** 1 year, 7 months ago

Never seen a score so high.

upvoted 4 times

🗨️ 👤 **adilkhan** Highly Voted 2 years, 2 months ago

both should be consistent prefix

upvoted 18 times

🗨️ 👤 **Makonee** 1 year, 11 months ago

Wrong. The Given answer is correct. Only Read Consistency can be set by client. Write stays at server setting which is session by default.

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level>

upvoted 21 times

🗨️ 👤 **Vladimir\_Gajinov** 1 year, 10 months ago

You are right, the given answer is correct.

Microsoft documentation states:

Overriding the default consistency level only applies to **reads** within the SDK client. **An account configured for strong consistency by default will still write and replicate data synchronously to every region in the account.** When the SDK client instance or request overrides this with Session or weaker consistency, reads will be performed using a single replica.

upvoted 1 times

🗨️ 👤 **giuliohome** 12 months ago

Still very confusing explanation. The other thing you must consider is that the level can only be relaxed. In conclusion I guess you have to configur prefix for write because it cannot be change while you can configure strong for read because it can be relaxed on a per SDK request basis.

upvoted 1 times

🗨️ 👤 **b1992** 1 year, 11 months ago

could you explain why?

upvoted 2 times

🗨️ 👤 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

🗨️ 👤 **Jay456** 9 months ago

Read operations will result in consistent prefix. Write operations will remain session because of the settings of the storage account.

upvoted 2 times

🗨️ 👤 **ciamp** 10 months, 2 weeks ago

For Azure Cosmos DB, when you provision an account with a default consistency level, you can override this setting on a per-request basis. Here's how you should configure the consistency levels for read and write operations to achieve a consistent prefix:

Read operations: Consistent Prefix

Write operations: Consistent Prefix

By setting both read and write operations to the Consistent Prefix level, you ensure that the data is read and written in the order in which it was updated, without requiring a strong consistency level, which can be more resource-intensive. This guarantees that reads never see out-of-order writes.

upvoted 1 times

🗨️ **Christian\_garcia\_martin** 11 months ago

Both are consistent prefix .

upvoted 1 times

🗨️ **giuliohome** 12 months ago

My answer would be that the resulting levels are both consistent prefix because it is written that they are set like that at account level and it is only possible to relax them via SDK (and only for read), so it would exclude strong and session as resulting levels for read or any other for write

upvoted 1 times

🗨️ **1989jay** 1 year, 5 months ago

This question was on exam. 09/01/2024. Passed 872.

upvoted 4 times

🗨️ **CarlosTheBoldest** 1 year, 6 months ago

I got this question on my exam, 2023Dec, go with what I remember was the most voted answer. Score 902, most of the questions were here, slightly different on wording because the Azure AD <-> Entra ID change. Case was City Power & Light. Good luck!

Important tip, you have access to Microsoft Learn during the exam!

upvoted 7 times

🗨️ **LittleBean** 1 year, 6 months ago

Hi, does all question were from exam topic? Did you bought contributor access? And what does it mean by case study city power and light. In free version I haven't come across this term/ former case study

upvoted 1 times

🗨️ **richardc7** 1 year, 9 months ago

got in 28/09/2023

upvoted 1 times

🗨️ **seanzhenxiao** 1 year, 10 months ago

Sorry, my first comment is error

I think should be :

read: consistent prefix,

write: session

1. default level is session

2. override only work for read

3. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account

strong > boundary > session > consistent prefix > eventually

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level>

upvoted 6 times

🗨️ **seanzhenxiao** 1 year, 10 months ago

I think should be session, session

1. default level is session

2. override only work for read

3. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level>

upvoted 1 times

🗨️ **Dianahu** 1 year, 11 months ago

answer is here: <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level>

Overriding the default consistency level only applies to reads within the SDK client.

default consistency is session

upvoted 6 times

  **[Removed]** 1 year, 11 months ago



Now I get it! Appreciated.

upvoted 2 times

  **Mark16** 2 years ago

stupid question xD

upvoted 4 times

  **dy0917** 2 years, 2 months ago

Any explanation?

For read operations with the requirement of consistent prefix, the resulting consistency level option is C. Consistent prefix.

Following is what ChatGPT said, but I am not convinced.

For write operations, there is no consistency level option called "Consistent prefix". The available options are Strong, Bounded staleness, Session, and Eventual. If you want to ensure consistency prefix for write operations, you should choose Strong consistency level.

upvoted 2 times

## DRAG DROP

You are developing an application to store millions of images in Azure blob storage. The images are uploaded to an Azure blob storage container named companyimages contained in an Azure blob storage account named companymedia. The stored images are uploaded with multiple blob index tags across multiple blobs in the container.

You must find all blobs whose tags match a search expression in the container. The search expression must evaluate an index tag named status with a value of final.

You need to construct the GET method request URI.

How should you complete the URI? To answer, drag the appropriate parameters to the correct request URI targets. Each parameter may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

## Parameters


Status='Final'  
Status <='Final'  
companymedia  
companyimages

## Answer Area


https://  .blob.core.windows.net/  ?restype=container&comp=blobs&where=

## Suggested Answer:


https://  companymedia .blob.core.windows.net/  companyimages ?restype=container&comp=blobs&where=  Status='Final'

 **adilkhan** Highly Voted 2 years, 2 months ago

Answer is correct!  
upvoted 26 times

 **DonH** Highly Voted 2 years ago

Just for information: I just had this question on my AZ204 exam - 16-jun-2023.  
I barely made it (with only 767 points) so I can't inform anyone if this answer is correct or not, just stating that this is an actual exam question.  
upvoted 7 times

 **NavinD** 1 year, 11 months ago

Do you have contributor access? what other websites did you study from?  
upvoted 1 times

 **MasonZ** 10 months, 3 weeks ago

Brother you won't see this question without contributor access now.  
upvoted 1 times

 **Tej\_007** 9 months, 3 weeks ago

previously around 200 questions were free. but now only 100 questions are free.  
upvoted 1 times

 **kysmf0100101** Most Recent 1 year, 5 months ago

no its not  
upvoted 1 times

 **Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850  
upvoted 2 times

 **lyggwtwtczxnhznebw** 1 year, 11 months ago

Correct  
upvoted 2 times

  **WorkFromHome** 2 years ago

Got this question on the exam at 2023/05/31

upvoted 3 times



HOTSPOT

-

You develop two Python scripts to process data.

The Python scripts must be deployed to two, separate Linux containers running in an Azure Container Instance container group. The containers must access external data by using the Server Message Block (SMB) protocol. Containers in the container group must run only once.

You need to configure the Azure Container Instance.

Which configuration value should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

Configuration Setting	Configuration Value
External data volume	<div>▼ Secret Empty directory Cloned git repo Azure file share</div>
Container restart policy	<div>▼ Never Always OnFailure</div>

### Answer Area

Suggested Answer:

Configuration Setting	Configuration Value
External data volume	<div>▼ Secret Empty directory Cloned git repo <b>Azure file share</b></div>
Container restart policy	<div>▼ <b>Never</b> Always OnFailure</div>

 **junix\_** Highly Voted 1 year, 5 months ago

I believe the answer is correct.

Azure file share: <https://learn.microsoft.com/en-us/azure/storage/files/files-smb-protocol?tabs=azure-portal>  
upvoted 17 times

 **ay\_m** Highly Voted 1 year, 4 months ago

Azure File Share is the only option that supports SMB  
upvoted 10 times

 **Vladimir\_Gajinov** 1 year, 4 months ago

You are right

upvoted 1 times

  **wafa\_chaari** Most Recent 5 months ago

correct

The SMB protocol is a network file sharing protocol that allows applications on a computer to read and write to files. Azure Files offers two industry-standard protocols for mounting Azure file share: the Server Message Block (SMB) protocol and the Network File System (NFS) protocol. Azure Files enables you to pick the file system protocol that is the best fit for your workload.

upvoted 2 times

  **1CY1** 6 months, 4 weeks ago

<https://learn.microsoft.com/en-us/azure/container-instances/container-instances-restart-policy>

Container restart policy : Never

Containers in the container group are never restarted. The containers run at most once.

upvoted 1 times

  **bgbgvfvf** 1 year ago

I think given answer is correct.

upvoted 1 times

  **arpannaprap** 1 year, 5 months ago

Answer is correct.



upvoted 4 times

  **[Removed]** 1 year, 5 months ago

Incorrect



Correct Answer: Empty Directory, Never

upvoted 1 times

  **Rithikk** 1 year, 1 month ago



Whyyyy????????????????

upvoted 2 times

  **Papee** 1 year, 4 months ago

why Empty Directory ?

upvoted 2 times

  **Ciupaz** 1 year, 1 month ago

Is not Empty Directory, but Azure File Share.

upvoted 2 times

  **yusuf\_eb** 1 year, 1 month ago

this guy is wrong

upvoted 1 times

## HOTSPOT

-

You are developing a static website hosted on Azure Blob Storage. You create a storage account and enable static website hosting.

The website must support the following requirements:

- Custom domain name
- Custom header values for all responses
- Custom SSL certificate

You need to implement the static website.

What should you configure? To answer, select the appropriate values in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area****Requirement****Configuration Value**

Custom domain name

▼

Blob index tags

Azure Content Delivery Network (CDN)

Cross-Origin Resource Sharing (CORS)

Azure Storage Service Encryption (SSE)

Custom header values

▼

Blob index tags

Azure Content Delivery Network (CDN)

Cross-Origin Resource Sharing (CORS)

Azure Storage Service Encryption (SSE)

Custom SSL certificate

▼

Blob index tags

Azure Content Delivery Network (CDN)

Cross-Origin Resource Sharing (CORS)

Azure Storage Service Encryption (SSE)

## Answer Area

### Requirement

### Configuration Value

Suggested Answer:

Custom domain name

▼

Blob index tags

Azure Content Delivery Network (CDN)

Cross-Origin Resource Sharing (CORS)

Azure Storage Service Encryption (SSE)

Custom header values

▼

Blob index tags

Azure Content Delivery Network (CDN)

Cross-Origin Resource Sharing (CORS)

Azure Storage Service Encryption (SSE)

Custom SSL certificate

▼

Blob index tags

Azure Content Delivery Network (CDN)

Cross-Origin Resource Sharing (CORS)

Azure Storage Service Encryption (SSE)


 **AbidooKing** Highly Voted 1 year, 5 months ago

Answer seems correct.

"Static websites have some limitations. For example, If you want to configure headers, you'll have to use Azure Content Delivery Network (Azure CDN)" and "To enable HTTPS, you'll have to use Azure CDN because Azure Storage doesn't yet natively support HTTPS with custom domains"


<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-static-website>

upvoted 21 times

 **Vukile** Highly Voted 1 year, 1 month ago

On exam 9 Nov 2023, went with given answer, score 865. Case Study: Farmers and Distributors

upvoted 6 times

 **Ciupaz** 1 year, 1 month ago

"Farmers and Distributor" is a new one (not present here on ExamTopics).

upvoted 4 times

 **priomkhan** 12 months ago

I am also looking for this case study.

upvoted 2 times

 **dom271219** 11 months, 2 weeks ago

it's in there : Munson's Pickles and preserves farms Topic 4, Q69, T5 ..., T6...,

upvoted 3 times

 **neelkanths** Most Recent 8 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

answer is correct...

upvoted 5 times

 **ericguo2020** 1 year, 3 months ago

Using Azure CDN to Specify Custom HTTP Headers for an Azure Static Website Hosted SPA: <https://medium.com/datadigest/using-azure-cdn-to-specify-custom-http-headers-for-an-azure-static-website-hosted-spa-41a9b9ec1674>


upvoted 2 times

 **Vladimir\_Gajinov** 1 year, 4 months ago

The answers are correct.

The most appropriate answer is Azure CDN, as it can fulfill all three requirements: custom domain name, custom header values, and a custom SSL certificate.

upvoted 3 times

 **Papee** 1 year, 4 months ago

### 1. Custom Domain Name:

- Configuration: Azure Content Delivery Network (CDN)

Azure CDN can be configured to map your custom domain name (e.g., [www.yourdomain.com](http://www.yourdomain.com)) to the Azure Blob Storage static website endpoint. This allows you to access your static website using your custom domain.

### 2. Custom Header Values for All Responses:

- Configuration: Blob index tag

You can use blob index tags to set custom header values for all responses from your static website. Blob index tags allow you to define metadata at the container or blob level that can be used as headers for the website content.

### 3. Custom SSL Certificate:

- Configuration: Azure Content Delivery Network (CDN)

To use a custom SSL certificate for your custom domain, you can configure Azure CDN to use the custom SSL certificate. This will enable secure communication between the user's browser and the CDN endpoint serving your static website content.

upvoted 3 times

  **NightshadeRC** 1 year, 5 months ago

Had this question in today's exam: 2023-07-26

upvoted 2 times

You are developing an inventory tracking solution. The solution includes an Azure Function app containing multiple functions triggered by Azure Cosmos DB. You plan to deploy the solution to multiple Azure regions.

The solution must meet the following requirements:

- Item results from Azure Cosmos DB must return the most recent committed version of an item.
- Items written to Azure Cosmos DB must provide ordering guarantees.

You need to configure the consistency level for the Azure Cosmos DB deployments.

Which consistency level should you use?

- A. consistent prefix
- B. eventual
- C. bounded staleness
- D. strong
- E. session

**Suggested Answer: D**

Community vote distribution

D (82%)

C (18%)

🗳️ 👤 **LuchoGonzalez** Highly Voted 1 year, 11 months ago

Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Source: <https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

upvoted 13 times

🗳️ 👤 **Xardas** Highly Voted 1 year, 11 months ago

I believe its C 'bounded staleness'

Because strong consistency is not supported for multiple regions

upvoted 11 times

🗳️ 👤 **manoptydakon** 1 year, 5 months ago

Yes, the Strong consistency level is supported in Azure Cosmos DB for multiple regions. When you configure Azure Cosmos DB with multiple regions, you have the option to choose the consistency level that fits your application's requirements. The Strong consistency level ensures that, regardless of the region being accessed, each read operation returns the most recent committed version of an item and maintains ordering guarantees. This is achieved by synchronously replicating writes across all regions before acknowledging the write operation.

upvoted 4 times

🗳️ 👤 **tarek0811** 1 year, 11 months ago

Azure Cosmos DB accounts configured with multiple write regions can't be configured for strong consistency as it isn't possible for a distributed system to provide an RPO of zero and an RTO of zero

<https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

upvoted 6 times

🗳️ 👤 **1CY1** 1 year ago

...strong consistency completes an operation only after ensuring that it has been committed to all regions within an account.

<https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels#write-latency-and-strong-consistency>

upvoted 2 times

🗳️ 👤 **oskx2** 1 year, 3 months ago

That is true but there is no mention of "multiple write regions" mentioned in the question.

You can have strong consistency and multiple regions (notice it is not multiple write regions): <https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels#write-latency-and-strong-consistency>

upvoted 5 times

🗳️ 👤 **vizay** Most Recent 2 months ago

Selected Answer: D

But it doesn't explicitly say that Cosmos DB itself is configured for multi-region writes or reads — only that the overall solution is deployed to multiple regions.

👉 So, does that mean Cosmos DB is multi-region?

Not necessarily. Just because the app is multi-region doesn't mean Cosmos DB is. You can still have a single-region Cosmos DB backend and multiple frontends (Function Apps) pointing to it.

upvoted 1 times

🗳️ 👤 **cmmr** 7 months, 1 week ago

This page explains it pretty well: <https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels#strong-consistency>

It shall be strong, as the write is only committed once it's replicated to all regions, and therefore all reading parties (from all regions) read the last inserted value (a tradeoff between availability + throughput and consistency)

upvoted 1 times

🗳️ 👤 **Vichu\_1607** 7 months, 3 weeks ago

Selected Answer: D

D. strong: Strong consistency guarantees that reads always return the most recent committed version of an item. It also provides linearizability, which ensures that all operations are seen in a strict order. This meets both requirements of returning the most recent committed version and providing ordering guarantees.

upvoted 1 times

🗳️ 👤 **0cc50bf** 10 months, 3 weeks ago

There is a trick here - the SOLUTION is deployed to multiple regions. That only contains the function app. Nowhere does it say that the COSMOS db is in multiple regions. Therefore, Strong is perfectly acceptable.

upvoted 3 times

🗳️ 👤 **Christian\_garcia\_martin** 11 months, 1 week ago

D is correct .According to the requirements that item results from Azure Cosmos DB must return the most recent committed version of an item and items written to Azure Cosmos DB must provide ordering guarantees, the Strong consistency level should be used.

Strong consistency offers linearizability, which is the highest consistency level available. It guarantees that once an operation completes, the result will be immediately visible to all subsequent operations. This ensures that the most recent committed version of an item is always returned and also provides ordering guarantees, but this comes with a trade-off as strong consistency may introduce higher latencies and reduced availability compared to other consistency levels.

upvoted 1 times

🗳️ 👤 **Munwalinwali** 12 months ago

Selected Answer: D

"Most recent"

upvoted 1 times

🗳️ 👤 **jobolesonihal** 1 year, 3 months ago

Selected Answer: D

The problem doesn't specify any condition (versions or time lag) which may allow bounded staleness.

Only Strong consistency guarantees the reading the more recent commit.

upvoted 3 times

🗳️ 👤 **FeriAZ** 1 year, 4 months ago

Selected Answer: D

Item results from Azure Cosmos DB must return the most recent committed version of an item: Strong consistency ensures that once a write operation is acknowledged, all subsequent read operations will return the most recent committed version of the data. This means that any read operation will see the effects of the latest write operation.

Items written to Azure Cosmos DB must provide ordering guarantees: Strong consistency also ensures ordering guarantees. It means that if one write operation is acknowledged before another, then the order of these operations is preserved for all subsequent read operations.

upvoted 1 times

🗳️ 👤 **Iednari** 1 year, 6 months ago

**Selected Answer: D**

Strong - doesn't mention multi-region writes - the multi-region functions are triggered hence reads  
upvoted 3 times

🗳️ 👤 **CarlosTheBoldest** 1 year, 7 months ago

**Selected Answer: C**

The key is the multiple regions  
upvoted 2 times

🗳️ 👤 **Jarynpl** 1 year, 7 months ago

**Selected Answer: C**

Azure Cosmos DB accounts configured with multiple write regions can't be configured for strong consistency as it isn't possible for a distributed system to provide an RPO of zero and an RTO of zero. Additionally, there are no write latency benefits on using strong consistency with multiple write regions because a write to any region must be replicated and committed to all configured regions within the account. This scenario results in the same write latency as a single write region account.

<https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels#strong-consistency-and-multiple-write-regions>

upvoted 1 times

🗳️ 👤 **Weam** 1 year, 7 months ago

but the question didn't mention that there will be multiple WRITE regions ?

upvoted 3 times

🗳️ 👤 **RikinPatel** 1 year, 9 months ago

**Selected Answer: D**

Strong is correct ans.:

Strong consistency: Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item. A client never sees an uncommitted or partial write. Users are always guaranteed to read the latest committed write.

Bounded staleness consistency:

With Bounded Staleness consistency, reads issued against a non-primary region may not necessarily return the most recent version of the data globally, but are guaranteed to return the most recent version of the data in that region, which will be within the maximum staleness boundary globally.

<https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

upvoted 3 times

🗳️ 👤 **tamerm** 1 year, 9 months ago

**Selected Answer: D**

I think it's D.

Strong consistency is not supported for multiple WRITE regions, and I don't see any such requirements in the question.

<https://learn.microsoft.com/en-us/azure/cosmos-db/high-availability#multiple-write-regions>

upvoted 2 times

🗳️ 👤 **RikinPatel** 1 year, 9 months ago

as per Que.: You plan to deploy the solution to multiple Azure regions.

upvoted 1 times

🗳️ 👤 **Vladimir\_Gajinov** 1 year, 10 months ago

**Selected Answer: C**

I think it's Bounded Staleness because it's not mentioned that read is executed right after write operation so I think it's acceptable to have small latency between write and read operations. I would go with this!

upvoted 1 times

🗳️ 👤 **drkpnh** 1 year, 10 months ago

But for Strong consistency, it also says Bounded Staleness in a multi-write account is an anti-pattern. So I really don't understand what the correct answer is.

Bounded Staleness or Strong?

upvoted 3 times



HOTSPOT

-

You are developing an application that runs in several customer Azure Kubernetes Service clusters. Within each cluster, a pod runs that collects performance data to be analyzed later. A large amount of data is collected so saving latency must be minimized.

The performance data must be stored so that pod restarts do not impact the stored data. Write latency should be minimized.

You need to configure blob storage.

How should you complete the YAML configuration? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

apiVersion: storage.k8s.io/v1

kind:

metadata:

name: data-store

provisioner: kubernetes.io/

parameters:

skuName: Premium\_LRS

reclaimPolicy:

## Answer Area

apiVersion: storage.k8s.io/v1

kind: 

▼  
PodStorage  
StorageClass  
PersistentVolume  
PersistentVolumeClaim

metadata:

name: data-store

provisioner: kubernetes.io/ 

▼  
azure-disk  
azure-file  
portworx-volume  
scaleio

parameters:

skuName: Premium\_LRS

reclaimPolicy: 

▼  
local  
retain  
delete

Suggested Answer:

🗄️ **richardc7** Highly Voted 1 year, 9 months ago

got in 28/09/23 , totally out of scope

upvoted 10 times

🗄️ **[Removed]** Highly Voted 1 year, 10 months ago

Sadly got this question 08/17/23 in the exam

upvoted 9 times

🗄️ **macobuzi** 1 year, 9 months ago

Thanks for mentioning, so this question is definitely not out of scope! In fact, I don't even believe "scope" is in Microsoft's dictionary

upvoted 16 times

🗄️ **StelOPapad4** 1 year, 1 month ago

Isn't Kubernetes a whole other Certification?

upvoted 2 times

🗄️ **8ac3742** Most Recent 10 months, 4 weeks ago

kind should be PersistentVolume and PersistentVolumeClaim to ensure data persistence.

provisioner should be kubernetes.io/azure-disk for Azure Disk.

skuName is set to Premium\_LRS for low latency.

reclaimPolicy is set to Retain to ensure data is not deleted when the pod is restarted.

upvoted 1 times

🗄️ **Christian\_garcia\_martin** 10 months, 2 weeks ago

in Kubernetes manifest you don't set the provisioner in a PV or PVC , you do in storageClass

upvoted 1 times

🗄️ **sekelsenmat** 1 year, 4 months ago

On exam 2024, went with given answer, score 872 or something. Case Study: Farmers and Distributors. I don't understand the question. The question states "configure blob", and then there is no option for blob configuration, but instead you configure Azure Files or Azure Disks?

And the worse of all, the configuration given is so obsolete that:

\*it was already removed from Kubernetes, as stated here: <https://kubernetes.io/docs/concepts/storage/storage-classes/#azure-disk>

\*it is only available in an example outside MS Learn (so unavailable during the exam)

Anyway, I'm still not sure what the right answer would be, but I went with azure-disk because it supports retentionPolicy, maybe it should be azure-files like some are saying, but its examples don't use retentionPolicy.

upvoted 4 times

🗨️ 👤 **lednari** 1 year, 6 months ago

StorageClass / azure-file / retain

upvoted 4 times

🗨️ 👤 **oCaio** 1 year, 10 months ago

Out of scope. Suggested answer seems correct. Ref:

<https://learn.microsoft.com/en-us/azure/aks/azure-csi-disk-storage-provision>

<https://kubernetes.io/docs/concepts/storage/storage-classes/#azure-disk>

Ref azure-file

<https://learn.microsoft.com/en-us/azure/aks/azure-csi-files-storage-provision#create-a-storage-class>

(Create a storage class)

<https://kubernetes.io/docs/concepts/storage/storage-classes/#azure-file>

Ref PodStorage / Persistent Volume / PersistentVolumeClaim

<https://kubernetes.io/docs/tasks/configure-pod-container/configure-persistent-volume-storage/>

Ref portworx-volume

<https://docs.portworx.com/portworx-enterprise/operations/operate-kubernetes/storage-operations/kubernetes-storage-101/volumes>

Ref Reclaim Policy

<https://kubernetes.io/docs/concepts/storage/storage-classes/#reclaim-policy>

upvoted 5 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

As someone mentioned to received this question recently, it is not out of scope. Actually, I am no longer care if it is out of scope because Microsoft doesn't give a f about it!

upvoted 4 times

🗨️ 👤 **Mpho7** 1 year, 11 months ago

Crazy that people are getting out-of-scope questions in the official exam. Perhaps its not counted in the final score. We wont know

upvoted 4 times

🗨️ 👤 **macobuzi** 1 year, 10 months ago

Or maybe, "out-of-scope" is not within Microsoft dictionary!

upvoted 6 times

🗨️ 👤 **BaoNguyen2411** 1 year, 11 months ago

Got this question 29/06/2023

upvoted 4 times

🗨️ 👤 **[Removed]** 1 year, 11 months ago

Out of scope! Nevertheless, seems to be correct.

upvoted 3 times

## HOTSPOT

-

## Case study

-

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

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

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

## To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

## Background

-

VanArsdel, Ltd. is a global office supply company. The company is based in Canada and has retail store locations across the world. The company is developing several cloud-based solutions to support their stores, distributors, suppliers, and delivery services.

## Current environment

-

## Corporate website

-

The company provides a public website located at <http://www.vanarsdelltd.com>. The website consists of a React JavaScript user interface, HTML, CSS, image assets, and several APIs hosted in Azure Functions.

## Retail Store Locations

-

The company supports thousands of store locations globally. Store locations send data every hour to an Azure Blob storage account to support inventory, purchasing and delivery services. Each record includes a location identifier and sales transaction information.

## Requirements

-

The application components must meet the following requirements:

## Corporate website

-

- Secure the website by using SSL.
- Minimize costs for data storage and hosting.
- Implement native GitHub workflows for continuous integration and continuous deployment (CI/CD).
- Distribute the website content globally for local use.
- Implement monitoring by using Application Insights and availability web tests including SSL certificate validity and custom header value verification.
- The website must have 99.95 percent uptime.

## Retail store locations

-

- Azure Functions must process data immediately when data is uploaded to Blob storage. Azure Functions must update Azure Cosmos DB by using native SQL language queries.
- Audit store sale transaction information nightly to validate data, process sales financials, and reconcile inventory.

## Delivery services

-

- Store service telemetry data in Azure Cosmos DB by using an Azure Function. Data must include an item id, the delivery vehicle license plate, vehicle package capacity, and current vehicle location coordinates.
- Store delivery driver profile information in Azure Active Directory (Azure AD) by using an Azure Function called from the corporate website.

## Inventory services

-

The company has contracted a third-party to develop an API for inventory processing that requires access to a specific blob within the retail store storage account for three months to include read-only access to the data.

## Security

-

- All Azure Functions must centralize management and distribution of configuration data for different environments and geographies, encrypted by using a company-provided RSA-HSM key.
- Authentication and authorization must use Azure AD and services must use managed identities where possible.

## Issues

-

## Retail Store Locations

-

- You must perform a point-in-time restoration of the retail store location data due to an unexpected and accidental deletion of data.
- Azure Cosmos DB queries from the Azure Function exhibit high Request Unit (RU) usage and contain multiple, complex queries that exhibit high point read latency for large items as the function app is scaling.

You need to implement the delivery service telemetry data.

How should you configure the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

Azure Cosmos DB

Value

API

▼

Core (SQL)  
Gremlin  
Table  
MongoDB

Partition Key

▼

Item id  
Vehicle license plate  
Vehicle package capacity  
Vehicle location coordinates

### Answer Area

Azure Cosmos DB

Value

API

▼

Core (SQL)  
Gremlin  
Table  
MongoDB

Partition Key

▼

Item id  
Vehicle license plate  
Vehicle package capacity  
Vehicle location coordinates

Suggested Answer:

**devex** Highly Voted 1 year, 10 months ago

API: "Core (SQL)"

Partition Key: "Item id"

See: <https://learn.microsoft.com/en-us/azure/cosmos-db/partitioning-overview#use-item-id-as-the-partition-key>

upvoted 31 times

**Jobalos009** 1 year ago

Great explanation ! Thanks for the link.

upvoted 1 times

**lyggwtwtczxnhznebw** 1 year, 10 months ago

That makes a lot more sense, thanks for the link

upvoted 1 times

**AndySmith** Highly Voted 1 year, 7 months ago

On exam 3-Nov-2023. Went with most-voted answer - 932/1000:

1) Core SQL

2) Item Id

Note: all 11 Qs from VanArsdel case study were on Exam.

upvoted 15 times

**Expandable** 1 year, 4 months ago

Why can I see only 1 question here  
is it because I dont have contributors access?  
upvoted 3 times

🗨️ 👤 **examtopicsLogin123** 6 months ago  
Further there will be more questions for case studies.  
Almost all of them appear after topic 6  
upvoted 1 times

🗨️ 👤 **vkalynenko** 1 year, 2 months ago  
I have this access, but I still see only one question...  
upvoted 2 times

🗨️ 👤 **shahid204** 1 year, 3 months ago  
yeah, even i just see 1 question, cases study has usually 10 questions  
upvoted 1 times

🗨️ 👤 **Karunan** Most Recent 1 year, 7 months ago  
Why its not License plate?  
upvoted 1 times

🗨️ 👤 **florianwicher** 5 months ago  
Two points that I noticed when reading <https://learn.microsoft.com/en-us/azure/cosmos-db/partitioning-overview#use-item-id-as-the-partition-key>

"Your partition key should be a property that (...) doesn't change."  
Surrogate keys never change by definition, a vehicles license plate might very well change.

ItemID also has a higher cardinality than Licence plate, because there are likely more items than vehicles.

The article also mentions highly-performing point reads when using the item ID.  
upvoted 2 times

🗨️ 👤 **Geirebass** 1 year, 10 months ago  
This should be Core(SQL) and Partition Key (item.id) i belive ?  
upvoted 3 times

🗨️ 👤 **WH16** 1 year, 10 months ago  
Why item.id is not used for Partition Key?  
upvoted 9 times

You create and publish a new Azure App Service web app.

User authentication and authorization must use Azure Active Directory (Azure AD).


You need to configure authentication and authorization.

What should you do first?

- A. Add an identity provider.
- B. Map an existing custom DNS name.
- C. Create and configure a new app setting.
- D. Add a private certificate.
- E. Create and configure a managed identity.

**Suggested Answer: A**

Community vote distribution



🗳️ 👤 **Firo** Highly Voted 1 year, 11 months ago

**Selected Answer: A**

A- Add an identity provide is correct

<https://learn.microsoft.com/en-us/azure/app-service/scenario-secure-app-authentication-app-service#3-configure-authentication-and-authorization>  
upvoted 15 times

🗳️ 👤 **Vladimir\_Gajinov** 1 year, 10 months ago

Correct - step 3

upvoted 1 times

🗳️ 👤 **lyggwtwtczxnhznebw** Highly Voted 1 year, 11 months ago

**Selected Answer: A**

Answer seems correct based on <https://learn.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad?tabs=workforce-tenant>  
upvoted 5 times

🗳️ 👤 **Munwalinwali** Most Recent 12 months ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/azure/app-service/scenario-secure-app-authentication-app-service#3-configure-authentication-and-authorization>  
upvoted 1 times

🗳️ 👤 **22312ce** 1 year, 4 months ago

Select Answer : E

upvoted 1 times

🗳️ 👤 **huuphong91** 1 year, 8 months ago

I think E is right

upvoted 2 times

🗳️ 👤 **nerojd** 1 year, 9 months ago

**Selected Answer: E**

Why not E. Create and configure a managed identity??

upvoted 2 times

🗳️ 👤 **Mattt** 7 months, 3 weeks ago

Managed identities allow the app to authenticate with other Azure services but are not necessary for configuring Azure AD authentication for users.

upvoted 1 times

🗳️ 👤 **cldmcp** 1 year, 8 months ago



I think it is because you need to implement user authentication and authorization. Managed identity is used to authenticate Azure resources

upvoted 8 times

## DRAG DROP

-

You have an Azure Cosmos DB for NoSQL account.

You plan to develop two apps named App1 and App2 that will use the change feed functionality to track changes to containers. App1 will use the pull model and App2 will use the push model.

You need to choose the method to track the most recently processed change in App1 and App2.

Which component should you use? To answer, drag the appropriate components to the correct apps. Each component may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Components	Answer Area	
	App	Component
Lease container		
Integrated cache	App1	
Continuation token	App2	

Answer Area		
	App	Component
Suggested Answer:	App1	Continuation token
	App2	Lease container

 **oCaio** Highly Voted 1 year, 4 months ago  
Seems correct.

App2 Push Model - Lease container

"When reading from the Azure Cosmos DB change feed, we usually recommend using a push model because you won't need to worry about: ... Storing state for the last processed change. If you are reading from the change feed processor, state is automatically stored in a lease container.

Ref:

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/read-change-feed#reading-change-feed-with-a-push-model>  
upvoted 11 times

 **FeriAZ** Highly Voted 11 months, 2 weeks ago  
Correct.


App1 (Pull Model): This model involves explicitly querying the change feed and managing the state of what has been read.

Best Component: Continuation Token. The pull model typically relies on continuation tokens to keep track of where the last read operation ended and to resume from that point.

App2 (Push Model): In this model, the change feed processor pushes changes to the application, and the application logic processes these changes.

Best Component: Lease Container. The push model, especially when using the change feed processor, leverages a lease container to maintain state and ensure reliable processing.

upvoted 9 times

 **wafa\_chaari** Most Recent 5 months ago  
correct :

You need to choose the method to track the most recently processed change in App1 and App2.

when reading from the Azure Cosmos DB change feed, we usually recommend using a push model because you won't need to worry about Storing state for the last processed change. If you are reading from the change feed processor, state is automatically stored in a lease container.

--

The change feed pull model allows you to consume the change feed at your own pace. Changes must be requested by the client and there's no automatic polling for changes. If you want to permanently "bookmark" the last processed change (similar to the push model's lease container), you'll need to save a continuation token.

upvoted 1 times



  **thecrash** 1 year, 2 months ago

Correct!

The change feed pull model allows you to consume the change feed at your own pace. Changes must be requested by the client and there's no automatic polling for changes. If you want to permanently "bookmark" the last processed change (similar to the push model's lease container), you'll need to save a continuation token.

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/read-change-feed>

upvoted 6 times

  **Ayaz\_b** 1 year, 3 months ago

Correct

upvoted 1 times

  **mario\_no\_code** 1 year, 5 months ago

The answer is correct. Link: <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-pull-model?tabs=dotnet#compare-to-the-change-feed-processor>

upvoted 4 times

You have a Linux container-based console application that uploads image files from customer sites all over the world. A back-end system that runs on Azure virtual machines processes the images by using the Azure Blobs API.

You are not permitted to make changes to the application.

Some customer sites only have phone-based internet connections.

You need to configure the console application to access the images.

What should you use?

- A. Azure BlobFuse
- B. Azure Disks
- C. Azure Storage Network File System (NFS) 3.0 support
- D. Azure Files

**Suggested Answer: C**

Community vote distribution

A (100%)

🗳️ 👤 **130nk3r5** Highly Voted 1 year, 6 months ago

**Selected Answer: A**

The correct answer is:

A. Azure BlobFuse

Azure BlobFuse allows you to access Azure Blob Storage from Linux and Azure services as if it were a local file system, without changing the application. This is particularly useful for your scenario where you can't modify the application and need to process images stored in Azure Blob Storage. BlobFuse provides the necessary interface between the application and Azure Blob Storage.

upvoted 8 times

🗳️ 👤 **Jass1nonly** 1 year, 4 months ago

That's a chatGPT answer?

upvoted 3 times

🗳️ 👤 **SachinV** Highly Voted 1 year, 7 months ago

Azure BlobFuse, as

Azure Disks are block storage devices that are not designed to be accessed directly by applications.

Azure Storage NFS 3.0 support is still in preview and is not yet available for all accounts.

Azure Files is a managed file share service that is not designed to be used with slow and unreliable internet connections.

upvoted 5 times

🗳️ 👤 **examtopicsLogin123** 5 months, 2 weeks ago

I agree with the answer, but don't quite agree with the explanation

Azure Disks - designed to work with Virtual Machines

Azure Container Storage - not in the list of answer options

<https://learn.microsoft.com/en-us/azure/storage/common/storage-introduction>

Azure Files - like a regular disk in a computer, I think this option should work

<https://learn.microsoft.com/en-us/azure/storage/files/storage-files-introduction>

Network File System (NFS) 3.0 protocol support for Azure Blob Storage - also seems like a good option. The article doesn't say anything about the "preview"

<https://learn.microsoft.com/en-us/azure/storage/blobs/network-file-system-protocol-support>

Azure BlobFuse - this seems to be the correct answer

<https://learn.microsoft.com/en-us/azure/storage/blobs/blobfuse2-what-is>

upvoted 2 times

🗲️ 👤 **f09229d** Most Recent 3 months ago

**Selected Answer: A**

In exam 30/3/2025 with this answer, I scored 907/1000

upvoted 2 times

🗲️ 👤 **Vichu\_1607** 7 months, 3 weeks ago

**Selected Answer: A**

A. Azure BlobFuse: BlobFuse is a virtual file system driver for Azure Blob storage. It allows you to mount a Blob storage container as a file system on Linux. This would enable the console application to access the images stored in Azure Blob storage without requiring any changes to the application itself.

upvoted 3 times

🗲️ 👤 **Luxlux** 10 months ago

**Selected Answer: A**

Azure BlobFuse is an open-source virtual file system driver for Azure Blob storage. It allows you to mount a Blob storage container as a file system on Linux.

Advantages:

Seamless Integration: Since BlobFuse allows you to mount Blob storage as a file system, your application can interact with the Blob storage as if it were a local file system. This means no changes are needed in your application code.

Performance: BlobFuse is optimized for performance and can handle large files and high throughput.

Compatibility: It works well with Linux-based systems and containerized applications.

Network Efficiency: It can handle intermittent connectivity and optimize data transfer over limited bandwidth connections, which is beneficial for sites with phone-based internet connections.

upvoted 3 times

🗲️ 👤 **Hendrikdb** 1 year, 4 months ago

Azure files is the selected answers according to google bbard and chatgpt.

upvoted 1 times

🗲️ 👤 **kysmf0100101** 1 year, 5 months ago

Linux is unbiased. we should ff

upvoted 2 times

🗲️ 👤 **Sakhana88** 1 year, 6 months ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/azure/storage/blobs/blobfuse2-what-is>

upvoted 5 times

🗲️ 👤 **Razvan\_C** 1 year, 7 months ago

Is this still a valid question for AZ-204? I don't see those resources in study guide.

upvoted 4 times

🗲️ 👤 **DaveOla** 1 year, 3 months ago

Microsoft is just crazy

upvoted 2 times

🗲️ 👤 **junkz** 1 year, 7 months ago

not completely clear on what "access the images" really means, but seeing that the application in VM uses blob API, stands to reason that the console application will interact with blob storage, not files. given that the console application is uploading images, it reasonably means that uploads via blobfuse, as this is a virtual file system driver for azure blob storage

upvoted 1 times

🗲️ 👤 **Ciupaz** 1 year, 7 months ago

**Selected Answer: A**

Azure BlobFuse is the correct answer in this scenario.

upvoted 3 times

🗲️ 👤 **dasunl** 1 year, 7 months ago

Hi what is this?

upvoted 1 times

## DRAG DROP

-

You are developing several microservices named serviceA, serviceB, and serviceC. You deploy the microservices to a new Azure Container Apps environment.

You have the following requirements:

- The microservices must persist data to storage.
- serviceA must persist data only visible to the current container and the storage must be restricted to the amount of disk space available in the container.
- serviceB must persist data for the lifetime of the replica and allow multiple containers in the replica to mount the same storage location.
- serviceC must persist data beyond the lifetime of the replica while allowing multiple containers to access the storage and enable per object permissions.

You need to configure storage for each microservice.

Which storage type should you use? To answer, drag the appropriate storage types to the correct microservices. Each storage type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

## Storage types

Azure Blob Storage

Azure Files storage

Ephemeral volume

Container file system

## Answer Area

## Microservice

serviceA

serviceB

serviceC

## Storage type

## Answer Area

Suggested Answer:

## Microservice

serviceA

serviceB



serviceC

## Storage type

Container file system

Ephemeral volume

Azure Files storage

 **130nk3r5**  1 year, 6 months ago

Given these options, the storage types for each microservice should be:

serviceA: Container file system The container file system is local to the container and is only visible to the current container. It is restricted to the amount of disk space available in the container.

serviceB: Ephemeral Storage Ephemeral Storage is temporary storage that persists for the lifetime of the replica. It allows multiple containers in the replica to mount the same storage location.

serviceC: Azure Blob Storage Azure Blob Storage is a highly scalable and durable object storage service that persists data beyond the lifetime of the replica. It allows multiple containers to access the storage and supports per-object permissions.

upvoted 9 times

🗨️ 👤 **Christian\_garcia\_martin** Most Recent 11 months, 1 week ago

A ----> container file system.

B ----> Ephemeral Storage.

C : While Azure Blob Storage is a valid option for persisting data beyond the lifetime of a replica and allows access from multiple containers, it lacks a key feature mentioned in the requirement: it does not provide BUILT-IN functionality for mounting the blob storage AS FILE

SYSTEM ACCESSIBLE to the containers , Azure Files does offer , so C ----> Azure Files

upvoted 3 times

🗨️ 👤 **009bdaa** 1 year ago

Given answer is correct. <https://learn.microsoft.com/en-us/azure/container-apps/storage-mounts?tabs=smb&pivots=azure-cli>

upvoted 1 times

🗨️ 👤 **jobolesonihal** 1 year, 3 months ago

The options are tricky. Container File System is essentially an Ephemeral Storage.

Ephemeral Storage has both the options Container Scoped and Replica Scoped. :)

The given answer is the most appropriate choice.

upvoted 2 times

🗨️ 👤 **Jarynpl** 1 year, 7 months ago

Answer looks correct.

upvoted 4 times

🗨️ 👤 **Jarynpl** 1 year, 7 months ago

Sorry. I'm mean:

ServiceA: Container file system

ServiceB: Ephemeral volumes (it lives as long as replica lives)

ServiceC: Azure Blob Storage (it lives outside of replica. So it's beyond the lifetime. And we want to have blob instead of file)

upvoted 5 times

🗨️ 👤 **junkz** 1 year, 7 months ago

ACA does not mount to blob storage(<https://learn.microsoft.com/en-us/azure/container-apps/storage-mounts>), AKS i believe it does. so original answer is right

upvoted 6 times

🗨️ 👤 **AlbertoBT** 1 year, 7 months ago

C: Azure Files storage

Azure Blob Storage don't allow: per object permissions

upvoted 5 times

DRAG DROP

-

You are developing a web service that will run on Azure virtual machines that use Azure Storage. You configure all virtual machines to use managed identities.

You have the following requirements:

- Secret-based authentication mechanisms are not permitted for accessing an Azure Storage account.
- Must use only Azure Instance Metadata Service endpoints.

You need to write code to retrieve an access token to access Azure Storage. To answer, drag the appropriate code segments to the correct locations. Each code segment may be used once or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

#### Code segment 1

`http://localhost:50342/oauth2/token`

`http://169.254.169.254:50432/oauth2/token`

`http://localhost/metadata/identity/oauth2/token`

`http://169.254.169.254/metadata/identity/oauth2/token`

#### Code segment 2

`import xml.etree.ElementTree as e  
e.parse(response.read())`

`import csv  
csv.DictReader(response.read())`

`import yaml  
return yaml.load(response.read())`

`import json  
return json.loads(response.read())`

#### Answer Area

`import urllib.request`

`url = "`

Code segment 1

`queryString = ""`

`response = urllib.request.urlopen(url + queryString)`


Code segment 2

#### Suggested Answer:

##### Answer Area

```
import urllib.request
url = "http://169.254.169.254/metadata/identity/oauth2/token"
queryString = ""
response = urllib.request.urlopen(url + queryString)

import json
return json.loads(response.read())
```

 **gfanco** 2 months, 1 week ago

CORRECT

<https://learn.microsoft.com/en-us/entra/identity/managed-identities-azure-resources/how-to-use-vm-token>

upvoted 1 times

 **examtopicsLogin123** 6 months, 3 weeks ago

Answer is correct

<https://learn.microsoft.com/en-us/entra/identity/managed-identities-azure-resources/how-to-use-vm-token#get-a-token-using-http>

<https://learn.microsoft.com/en-us/azure/virtual-machines/instance-metadata-service?tabs=windows>

and there is the same question, but for C#

<https://www.examtopics.com/discussions/microsoft/view/49787-exam-az-204-topic-3-question-12-discussion/>

upvoted 2 times



🗨️ 👤 **abhishekgraphs** 1 year, 4 months ago

I asked Google AI tool the following question: How do I retrieve an access token using Azure Instance Metadata Service endpoints ?

I think the given answer here is correct!

upvoted 3 times

🗨️ 👤 **manopeydaKon** 1 year, 5 months ago

Not relevant question and answer!

upvoted 3 times

🗨️ 👤 **AbdullahFadhil** 11 months ago

It is part of the exam, you can check Microsoft website.

upvoted 1 times

HOTSPOT

-

You are developing an Azure Function app.

The Azure Function app must enable a WebHook to read an image from Azure Blob Storage and create a new Azure Cosmos DB document.

You need to implement the Azure Function app.

Which configuration should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

Trigger	Input binding	Output binding
<div>HTTP</div> <div>Timer</div> <div>Blob Storage</div> <div>Azure Cosmos DB</div>	<div>HTTP</div> <div>Timer</div> <div>Blob Storage</div> <div>Azure Cosmos DB</div>	<div>HTTP</div> <div>Timer</div> <div>Blob Storage</div> <div>Azure Cosmos DB</div>

### Answer Area

Suggested Answer:

Trigger	Input binding	Output binding
<div>HTTP</div> <div>Timer</div> <div>Blob Storage</div> <div>Azure Cosmos DB</div>	<div>HTTP</div> <div>Timer</div> <div>Blob Storage</div> <div>Azure Cosmos DB</div>	<div>HTTP</div> <div>Timer</div> <div>Blob Storage</div> <div>Azure Cosmos DB</div>

 Neckbread  11 months, 3 weeks ago

Given Answer is Correct:

A- HTTP Trigger for webhook

B- Input binding for blob storage to read image

C- Output binding to output a document to Cosmos DB

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-blob-input?tabs=python-v2%2Cisolated-process%2Cnodejs-v4&pivots=programming-language-csharp>

<https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-cosmosdb-v2-output?tabs=python-v2%2Cisolated-process%2Cnodejs-v4%2Cextensionv4&pivots=programming-language-csharp>

upvoted 9 times

🗲️ 👤 **sky\_may** Most Recent 11 months, 3 weeks ago

provided answer is correct: <https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-http-webhook-trigger?tabs=python-v2%2Cisolated-process%2Cnodejs-v4%2Cfunctions-v2&pivots=programming-language-csharp>

upvoted 4 times

🗲️ 👤 **AzDeveloper** 11 months, 3 weeks ago

Answer is correct because you can route storage events to webhook

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-quickstart?toc=%2Fazure%2Fevent-grid%2Ftoc.json>

upvoted 3 times

🗲️ 👤 **4cca361** 11 months, 4 weeks ago

Given answer is false:

Blob storage, Blob storage, Azure Cosmos DB

upvoted 4 times

🗲️ 👤 **laminall** 3 months, 3 weeks ago

WebHook -> HTTP

upvoted 1 times

You create an Azure Cosmos DB for NoSQL database.

You plan to use the Azure Cosmos DB .NET SDK v3 API for NoSQL to upload the following files:

File Name	File Size
File1	1MB
File2	2MB
File3	3MB
File4	4MB
File5	5MB

You receive the following error message when uploading the files: "413 Entity too large".

You need to determine which files you can upload to the Azure Cosmos DB for NoSQL database.

Which files can you upload?

- A. File1, File2, File3, File4, and File5
- B. File1 and File2 only
- C. File1, File2, and File3 only
- D. File1, File2, File3, and File4 only
- E. File1 only

**Suggested Answer: B**

Community vote distribution


B (95%)

5%

 **Jedi** Highly Voted 1 year, 5 months ago

**Selected Answer: B**

Correct - 2MB limit - <https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits#per-item-limits>  
upvoted 10 times

 **jobolesoniha1** Highly Voted 1 year, 3 months ago

**Selected Answer: B**

2 GB it is.

Maximum size of an item = 2 MB (UTF-8 length of JSON representation) <sup>1</sup>

<sup>1</sup> Large document sizes up to 16 MB are supported with Azure Cosmos DB for MongoDB only.

upvoted 5 times

 **f09229d** Most Recent 3 months ago

**Selected Answer: B**

In exam 30/3/2025 with this answer, I scored 907/1000

upvoted 2 times

 **MasonZ** 10 months, 2 weeks ago

**Selected Answer: B**

Select B.

"Azure Cosmos DB limits single request's size to 2MB."

<https://learn.microsoft.com/en-us/azure/data-factory/connector-azure-cosmos-db?tabs=data-factory#sink-transformation>

upvoted 1 times

 **Ushakanth** 1 year, 4 months ago

**Selected Answer: B**

Option B is correct. Each file has a maximum size of 2 MB. So File 1, File 2 Only can be uploaded.

upvoted 3 times

🗨️ **harbox** 1 year, 3 months ago

this correct: <https://learn.microsoft.com/en-us/answers/questions/692455/actual-document-size-limit-for-azure-cosmos-db-s-a>

While it is possible to store more than 2MB

upvoted 1 times

🗨️ **emysa341** 1 year, 5 months ago

**Selected Answer: B**

B is correct

upvoted 1 times

🗨️ **manopeydakon** 1 year, 5 months ago

E, because the total of files in the batch will exceed 2 MB

upvoted 1 times

🗨️ **sky\_may** 1 year, 5 months ago

Correct Answer is E: Azure Cosmos DB limits single request's size to 2MB. The formula is Request Size = Single Document Size \* Write Batch Size. If you hit error saying "Request size is too large.", reduce the writeBatchSize value in copy sink configuration

<https://learn.microsoft.com/en-us/azure/data-factory/connector-azure-cosmos-db?tabs=data-factory>

upvoted 2 times

🗨️ **AzDeveloper** 1 year, 5 months ago

**Selected Answer: E**

I think the correct answer is E not B because the total size for File1 and File2 will be 3 MB

upvoted 1 times

🗨️ **Woksi** 1 year, 4 months ago

The requirements do not specify that they MUST be uploaded in one batch

upvoted 2 times

🗨️ **1CY1** 12 months ago

I'll go with the agreed answer but it also did not say the files are uploaded individually.

So the answer E has some merit.

upvoted 1 times

You are developing an app to store globally distributed data in several Azure Blob Storage containers. Each container hosts multiple blobs where each instance of the app will store the data. You enable versioning and soft delete for the blobs.

App testing and incorrect code have frequently corrupted data. Development of the app must allow data to be restored to a previous day for testing.

You need to configure the storage account to support point-in-time restore.

What should you do?


- A. Enable the change feed on the storage account to begin capturing and recording changes.
- B. Configure object replication and specify replication rules.
- C. Create a snapshot of the blob in the hot tier.
- D. Configure an immutability policy that is scoped to a blob version.

**Suggested Answer: A**

  **Vichu\_1607** 7 months, 3 weeks ago

**Selected Answer: A**

A. Enable the change feed on the storage account to begin capturing and recording changes.  
upvoted 1 times

  **KeithSh** 8 months, 2 weeks ago

Looks like the answer is correct.

<https://learn.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-manage?tabs=portal#enable-and-configure-point-in-time-restore>  
upvoted 3 times

  **passme847** 8 months, 2 weeks ago

Prerequisites for point-in-time restore

Point-in-time restore requires that the following Azure Storage features be enabled before you can enable point-in-time restore:

Soft delete (already done)

Change feed (answer)

Blob versioning (already done)

<https://learn.microsoft.com/en-us/azure/storage/blobs/point-in-time-restore-overview>

upvoted 4 times

  **KeithSh** 8 months, 2 weeks ago

From the link above:

Enable and configure point-in-time restore

Before you enable and configure point-in-time restore, enable its prerequisites for the storage account: soft delete, change feed, and blob versioning. For more information about enabling each of these features, see these articles:

Enable soft delete for blobs

Enable and disable the change feed

Enable and manage blob versioning

upvoted 1 times



  **mahomes15** 8 months, 2 weeks ago

Given the need for point-in-time restore and the capabilities of the options:

Correct Answer: C. Create a snapshot of the blob in the hot tier.

Reasoning: Creating a snapshot allows you to capture the current state of the blob and restore it later, which meets the requirement of restoring to a previous day for testing.

upvoted 1 times

  **Mattt** 7 months, 3 weeks ago

It doesn't offer an automatic way to capture all changes or restore from a specific day across all blobs in the container.

upvoted 1 times

A company uses an Azure Blob Storage for archiving.

The company requires that data in the Blob Storage is only in the archive tier.

You need to ensure data copied to the Blob Storage is moved to the archive tier.



What should you do?

- A. Use a Put Block List operation with a request header of x-ms-immutability-policy-mode.
- B. Create a lifecycle policy with an action of tierToArchive and configure daysAfterModificationGreaterThan for 0.
- C. Use a Put Blob operation with a request header of x-ms-immutability-policy-until-date.
- D. Create a lifecycle policy with an action of tierToArchive and configure a filter for blobIndexMatch.

**Suggested Answer: B**

Community vote distribution

B (100%)

  **gfanco** 2 months, 1 week ago

**Selected Answer: B**

"


Name: daysAfterModificationGreaterThan

Description: Value indicating the age in days after last modification int

Value: Constraints: Min value = 0

"

<https://learn.microsoft.com/en-us/azure/templates/microsoft.storage/storageaccounts/managementpolicies?pivots=deployment-language-bicep>  
upvoted 1 times

  **gfanco** 2 months, 1 week ago

```
tierToArchive: {
  daysAfterCreationGreaterThan: int
  daysAfterLastAccessTimeGreaterThan: int
  daysAfterLastTierChangeGreaterThan: int
  daysAfterModificationGreaterThan: int <--
}
```

upvoted 1 times

  **cmmr** 7 months, 1 week ago

B would require at least 1 day to reach the threshold.

D would allow tagging blobs and immediately applying the tierToArchive policy.

I go with D

upvoted 1 times

  **cmmr** 7 months, 1 week ago

I stand corrected:

"Archive data after ingest


Some data stays idle in the cloud and is rarely, if ever, accessed. The following lifecycle policy is configured to archive data shortly after it's ingested. This example transitions block blobs in a container named archivecontainer into an archive tier. The transition is accomplished by acting on blobs 0 days after last modified time."



<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

It's B


upvoted 3 times

  **Vichu\_1607** 7 months, 3 weeks ago

**Selected Answer: B**

B. Create a lifecycle policy with an action of tierToArchive and configure daysAfterModificationGreaterThan for 0.

upvoted 2 times

  **overhill** 8 months ago

I think B is the right Answer

1) Creating a lifecycle policy with TierToArchive is necessary to move blobs to Archive

and setting daysAfterModificationGreaterThan will set the days after modification to 0.

From MSFT

daysAfterModificationGreaterThan:

Value indicating the age in days after last modification

<https://learn.microsoft.com/en-us/javascript/api/@azure/arm-storage/dateaftermodification?view=azure-node-latest>

upvoted 1 times

  **overhill** 8 months ago

First question I see without a single comment..

Well I also think B is the answer but I'm not sure at all

upvoted 1 times

## HOTSPOT

-

You have the following data lifecycle management policy:

```
{
  "rules": [
    {
      "enabled": true,
      "name": "Policy1",
      "type": "Lifecycle",
      "definition": {
        "actions": {
          "baseBlob": {
            "tierToArchive": {
              "daysAfterModificationGreaterThan": 0
            }
          }
        },
        "filters": {
          "blobIndexMatch": [
            {
              "name": "Customer",
              "op": "==",
              "value": "Adatum"
            }
          ]
        }
      }
    }
  ]
}
```

You plan to implement an Azure Blob Storage account and apply to it Policy1. The solution should maximize resiliency and performance.

You need to configure the account to support the policy.

Which redundancy option and storage account type should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

## Answer Area

### Storage account configuration

Configuration setting	Value
Redundancy	<div><div></div><div>ZRS</div><div>RA-GRS</div><div>RA-GZRS</div></div>
Account type	<div><div></div><div>premium page blob</div><div>premium block blob</div><div>general-purpose v2</div></div>

## Answer Area

### Storage account configuration

Suggested Answer:

Configuration setting	Value
Redundancy	<div><div></div><div>ZRS</div><div>RA-GRS</div><div>RA-GZRS</div></div>
Account type	<div><div></div><div>premium page blob</div><div>premium block blob</div><div>general-purpose v2</div></div>

 **KeithSh** Highly Voted 8 months, 2 weeks ago

Second answer incorrect. Should be General-purpose V2

<https://learn.microsoft.com/en-us/azure/storage/common/storage-redundancy#supported-storage-account-types>

upvoted 11 times

 **space\_cadet** 5 months, 3 weeks ago

Answers: RA-GRS and General-purpose V2

Premium block blob only has two redundancy options; LRS and ZRS

Archive tier is not supported on Zone redundant accounts

Links:

<https://learn.microsoft.com/en-us/azure/storage/common/storage-account-overview?source=recommendations#types-of-storage-accounts>

<https://learn.microsoft.com/en-us/azure/storage/common/storage-redundancy#supported-storage-account-types>

upvoted 4 times

 **wafa\_chaari** Most Recent 5 months ago

RA-GRS -->

The archive tier for Blob Storage isn't currently supported for ZRS, GZRS, or RA-GZRS accounts.

General-purpose V2 -->

<https://learn.microsoft.com/en-us/answers/questions/864000/difference-between-premium-vs-standard-general-pur>

upvoted 2 times

 **overhill** 8 months ago

My answers

-> RA-GZRS

-> General Purpose V2



Reason:

ZRS is just zone redundant in a single region

GRS is regionally redundant but doesn't allow reads

GZRS is regionally redundant and zone redundant (Maximum resilience)

upvoted 2 times

  **Tttttttaa** 6 months, 4 weeks ago

The archive tier for Blob Storage is currently supported for LRS, GRS, and RA-GRS accounts, but not for ZRS, GZRS, or RA-GZRS accounts

<https://learn.microsoft.com/en-us/azure/storage/common/storage-redundancy>

upvoted 4 times

  **overhill** 8 months ago



I mean:

ZRS is just zone redundant in a single region

(RA-)GRS is regionally redundant

(RA-)GZRS is regionally redundant and zone redundant (Maximum resilience)

upvoted 2 times

  **c75314a** 8 months, 2 weeks ago

"Microsoft recommends using GZRS for applications requiring maximum consistency, durability, and availability, excellent performance, and resilience for disaster recovery."

Only Standard general-purpose v2 (StorageV2)1 can handle GRS/RA-GRS

<https://learn.microsoft.com/en-us/azure/storage/common/storage-redundancy>

So, I would say RA-GRS with general purpose v2

upvoted 4 times

## HOTSPOT

-

You have an Azure Cosmos DB for NoSQL API account named account1. Multiple instances of an on-premises application named app1 read data from account1.

You plan to implement integrated cache for connections from the instances of app to account1.

You need to set the connection mode and maximum consistency level of app1.

Which values should you use for the configuration settings? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area****Connectivity mode and consistency level settings**

Configuration setting	Value
Connectivity mode	<div>direct mode gateway mode using standard gateway gateway mode using dedicated gateway</div>
Consistency level	<div>strong session eventual</div>

**Answer Area****Connectivity mode and consistency level settings****Suggested Answer:**


Configuration setting	Value
Connectivity mode	<div>direct mode gateway mode using standard gateway gateway mode using dedicated gateway</div>
Consistency level	<div>strong session eventual</div>

 **c75314a**  8 months, 2 weeks ago

Okay, so gateway mode using dedicated gateway has to be correct.

"The integrated cache uses the dedicated gateway within your Azure Cosmos DB account." (<https://learn.microsoft.com/en-us/azure/cosmos-db/integrated-cache>)

The second one I'm not sure about. "The integrated cache supports read requests with session and eventual consistency only." (<https://learn.microsoft.com/en-us/azure/cosmos-db/integrated-cache>). In my mind there are two correct answers here, but please correct me if I'm wrong. On that page they do talk a lot about session consistency and that leads me to believe that session is the correct answer  
upvoted 5 times

 **space\_cadet** 5 months, 3 weeks ago

session is also a stronger consistency level (maximum) when compared to eventual  
upvoted 1 times

 **overhill** 8 months ago

agreed

upvoted 1 times

  **Mattt** 7 months, 3 weeks ago

Consistency level should be session.

integrated cache only supports read requests with session and eventual consistency. Session consistency is recommended when the application has multiple instances or sessions and requires consistent reads per user session.

The se

upvoted 1 times

  **wafa\_chaari** Most Recent 5 months ago

seems correct

upvoted 1 times

  **overhill** 8 months ago

According to the documentation it seems to be correct

upvoted 1 times

You are developing a Cosmos DB solution that will be deployed to multiple Azure regions.

Your solution must meet the following requirements:

- Read operations will never receive write operations that are out of order.
- Maximize concurrency of read operations in all regions.

You need to choose the consistency level for the solution.

Which consistency level should you use?

- A. session
- B. eventual
- C. bounded staleness
- D. consistent prefix

**Suggested Answer: D**

*Community vote distribution*



**de55d65** 7 months, 1 week ago

**Selected Answer: D**

What we can rule out is B eventual, since there can be out-of-order reads, and C, since the question does not contain any information about a boundary.

If the wording of the question was "consistency" instead of "concurrency", I would go for A "session". Therefore I'd expect the lower the consistency, the higher the concurrency. I would go for D Consistent prefix.

Reference: <https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

No words about concurrency there.

upvoted 9 times

**patrykcpt** 3 weeks, 1 day ago

**Selected Answer: D**

Since consistent prefix allows for a consistent read for e.g. 2 write operations (documents aren't returned in mixed versions), then same applies for 3, 4, 5 etc., which means more write operations can be done in parallel. Hence the answer D, consistent prefix allows for bigger concurrency.

upvoted 1 times

**DanielRO** 5 months, 3 weeks ago

**Selected Answer: D**

ChatGPT response: Consistent prefix.

upvoted 1 times

**Jay456** 7 months, 1 week ago

**Selected Answer: D**

After reading the documentation I fully agree with de55d65's explanation and choose D as well.

upvoted 2 times

**Vichu\_1607** 7 months, 3 weeks ago

**Selected Answer: C**

C. bounded staleness: Bounded staleness consistency ensures that reads never see out-of-order writes and provides a lag window within which the data might be stale. This consistency level guarantees that reads are consistent within a specified staleness window, which meets the requirement of never receiving out-of-order writes. It also allows for high concurrency of read operations across multiple regions.

upvoted 3 times

**Mattt** 8 months, 1 week ago

Selected Answer: C

C is the correct answer

Bounded staleness guarantees that read operations will never see writes out of order, which directly addresses your requirement for ensuring that read operations never receive write operations that are out of order.

It also maximizes the concurrency of read operations in multiple regions while providing a guarantee that reads lag behind writes by a certain staleness window (defined in terms of time or number of versions). This balance between consistency and availability fits your scenario well.

upvoted 4 times

🗨️ 👤 **passme847** 8 months, 2 weeks ago

Selected Answer: D

"Consistent prefix consistency, the reads never see out of order writes for a transactional batch of writes"

<https://learn.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

upvoted 1 times

🗨️ 👤 **Mattt** 8 months, 1 week ago

But it doesn't maximize the concurrency

upvoted 2 times

🗨️ 👤 **Jay456** 7 months, 1 week ago

Neither does bounded staleness

upvoted 1 times

🗨️ 👤 **ns4098** 8 months, 2 weeks ago

Selected Answer: C

Why not C?

upvoted 2 times



You have an Azure Queue Storage named queue1.

You plan to develop code that will process messages in queue1.

You need to implement a queue operation to set the visibility timeout value of individual messages in queue1.

Which two operations can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Peek at a message in the queue.
- B. Delete a message in the queue.
- C. Add a message to the queue.
- D. Update a message in the queue.
- E. Receive a message from the queue.

**Suggested Answer:** CD

 **ns4098** Highly Voted 8 months, 2 weeks ago

To set the visibility timeout of individual messages in Azure Queue Storage, you can use the following operations:

D. Update a message in the queue

Explanation: The Update operation allows you to modify a message's content and its visibility timeout. You can use this operation to update the visibility timeout for a message that has already been retrieved from the queue.

E. Receive a message from the queue

Explanation: When you retrieve (or "receive") a message from the queue, you can specify the visibility timeout at that point. This is useful when you want to process the message but delay its visibility to other consumers.

upvoted 6 times

 **Vichu\_1607** Most Recent 7 months, 3 weeks ago

**Selected Answer:** DE

D. Update a message in the queue.

E. Receive a message from the queue.

upvoted 2 times

 **mahomes15** 8 months ago

**Selected Answer:** DE

E. Receive a message from the queue: When you receive a message, it becomes invisible to other consumers for the duration of the visibility timeout. After receiving it, you can modify its visibility timeout by using the Update operation or by changing the visibility timeout parameter when processing the message.

D. Update a message in the queue: This operation allows you to change the visibility timeout of a message that has been received. You can use this to reset the visibility timeout to a new value, effectively controlling how long the message remains hidden from other consumers.

upvoted 3 times

 **Mattt** 8 months, 1 week ago

**Selected Answer:** DE

D. Update a message in the queue: The UpdateMessage operation allows you to modify the content of a message as well as reset its visibility timeout value.

E. Receive a message from the queue: The ReceiveMessage operation retrieves the message from the queue and allows you to set a visibility timeout, which defines how long the message will be invisible to other consumers.

upvoted 3 times

## HOTSPOT

-

## Case study

-

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

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

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

## To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

## Background

-

Fourth Coffee is a global coffeehouse chain and coffee company recognized as one of the world's most influential coffee brands. The company is renowned for its specialty coffee beverages, including a wide range of espresso-based drinks, teas, and other beverages. Fourth Coffee operates thousands of stores worldwide.

## Current environment

-

The company is developing cloud-native applications hosted in Azure.

## Corporate website

-

The company hosts a public website located at <http://www.fourthcoffee.com/>. The website is used to place orders as well as view and update inventory items.

## Inventory items

-

In addition to its core coffee offerings, Fourth Coffee recently expanded its menu to include inventory items such as lunch items, snacks, and merchandise. Corporate team members constantly update inventory. Users can customize items. Corporate team members configure inventory items and associated images on the website.

## Orders

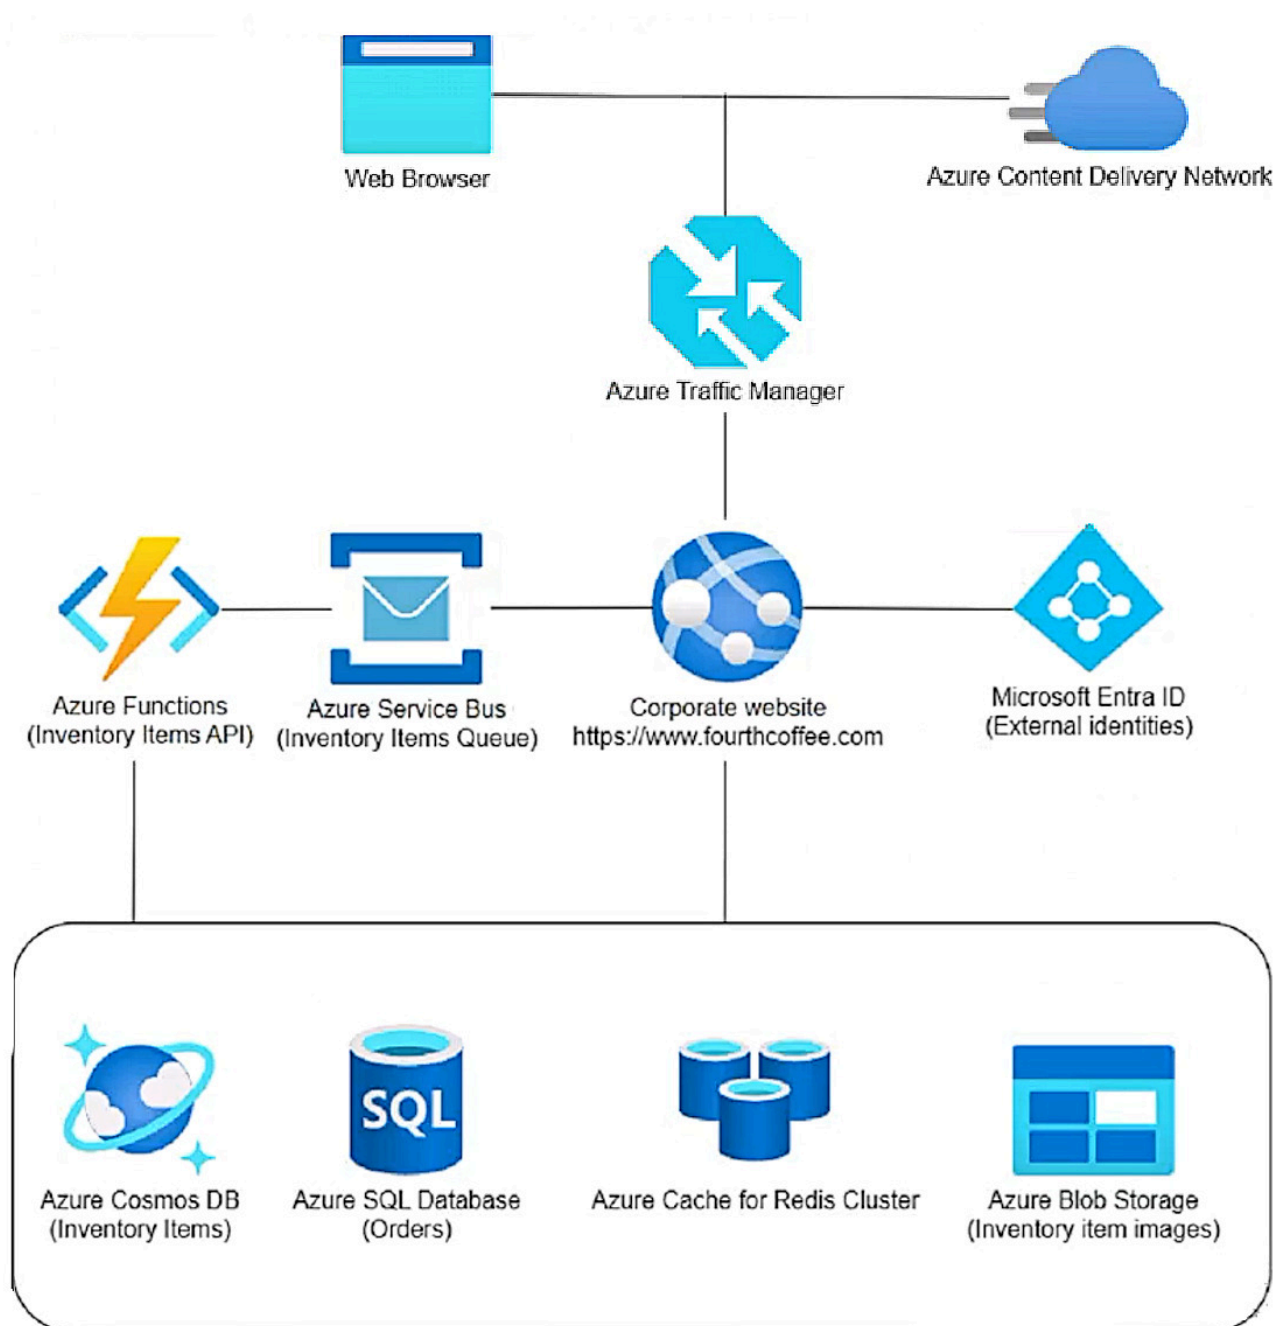
-

Associates in the store serve customized beverages and items to customers. Orders are placed on the website for pickup.

The application components process data as follows:

1. Azure Traffic Manager routes a user order request to the corporate website hosted in Azure App Service.
2. Azure Content Delivery Network serves static images and content to the user.
3. The user signs in to the application through a Microsoft Entra ID for customers tenant.
4. Users search for items and place an order on the website as item images are pulled from Azure Blob Storage.
5. Item customizations are placed in an Azure Service Bus queue message.
6. Azure Functions processes item customizations and saves the customized items to Azure Cosmos DB.
7. The website saves order details to Azure SQL Database.
8. SQL Database query results are cached in Azure Cache for Redis to improve performance.

The application consists of the following Azure services:



#### Requirements

-

The application components must meet the following requirements:

- Azure Cosmos DB development must use a native API that receives the latest updates and stores data in a document format.
- Costs must be minimized for all Azure services.
- Developers must test Azure Blob Storage integrations locally before deployment to Azure. Testing must support the latest versions of the Azure Storage APIs.

#### Corporate website

- 
- User authentication and authorization must allow one-time passcode sign-in methods and social identity providers (Google or Facebook).
- Static web content must be stored closest to end users to reduce network latency.

#### Inventory items

- 
- Customized items read from Azure Cosmos DB must maximize throughput while ensuring data is accurate for the current user on the website.
- Processing of inventory item updates must automatically scale and enable updates across an entire Azure Cosmos DB container.
- Inventory items must be processed in the order they were placed in the queue.
- Inventory item images must be stored as JPEG files in their native format to include exchangeable image file format (data) stored with the blob data upon upload of the image file.
- The Inventory Items API must securely access the Azure Cosmos DB data.

#### Orders

- 
- Orders must receive inventory item changes automatically after inventory items are updated or saved.

#### Issues

- 
- Developers are storing the Azure Cosmos DB credentials in an insecure clear text manner within the Inventory Items API code.
- Production Azure Cache for Redis maintenance has negatively affected application performance.

You need to save customized items to Azure Cosmos DB.

Which Azure Cosmos DB configuration should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

#### Azure Cosmos DB configuration

Configuration option	Configuration value
API	<div>▼</div> <div>NoSQL</div> <div>MongoDB</div> <div>PostgreSQL</div> <div>Apache Gremlin</div>
Consistency level	<div>▼</div> <div>Strong</div> <div>Session</div> <div>Eventual</div> <div>Bounded Staleness</div>

### Answer Area

#### Azure Cosmos DB configuration

Suggested Answer:

##### Configuration option

API

Consistency level

##### Configuration value

NoSQL

MongoDB

PostgreSQL

Apache Gremlin

Strong

Session

Eventual

Bounded Staleness

 **ns4098** Highly Voted 8 months, 2 weeks ago

1. API:

NoSQL

Reasoning: Since the development must use a native API that stores data in a document format, the NoSQL API is the correct choice. Azure Cosmos DB's NoSQL API stores data in JSON documents and is the most commonly used API for document-based use cases, which aligns with the storage of customized items in the scenario.

2. Consistency Level:

Session

Reasoning: The scenario requires that customized items maximize throughput while ensuring that the data is accurate for the current user. Session consistency provides the best balance between performance and data accuracy for scenarios where multiple reads and writes are performed by the same user session. It ensures that a user reads their own writes, which is important for ensuring that customized items are accurate for the current user.

Final Answer:

API: NoSQL

Consistency level: Session


upvoted 10 times

 **KeithSh** 8 months, 1 week ago

Agree first answer should be NoSQL. This is a new application so best option is NoSQL.

<https://learn.microsoft.com/en-us/azure/cosmos-db/choose-api#considerations-when-choosing-an-api>

upvoted 2 times

 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

 **examtopicsLogin123** 4 months, 2 weeks ago

To prevent reading the cases multiple times:

Please see the spots below where you can find the questions (topic/question/description)

Fourth Coffee

3/60 save customized items to Azure Cosmos DB

3/64 store inventory item images

4/75 secure the corporate website for users

4/77 securely access inventory items

5/56 mitigate the Azure Cache for Redis issue

5/57 serve static content from the corporate website

5/60 implement code to process inventory changes and update orders

6/53 implement a function by using Azure Functions to process customized items



6/54 implement the processing of enqueue inventory items

for the last two I provided the numbers specified on the page, but they have different numbers in URL and title:

<https://www.examttopics.com/discussions/microsoft/view/151609-exam-az-204-topic-6-question-63-discussion/>

<https://www.examttopics.com/discussions/microsoft/view/150590-exam-az-204-topic-6-question-62-discussion/>

upvoted 2 times

  **e66b406** 5 months, 2 weeks ago

NoSQL and Session

upvoted 1 times

## HOTSPOT

-

A company has an Azure Cosmos DB for NoSQL account. The account is configured for session consistency. Data is written to a single Azure region and data can be read from three Azure regions.

An application that will access the Azure Cosmos DB for NoSQL container data using an SDK has the following requirements:

- Reads from the application must return the most recent committed version of an item from any Azure region.
- The container items should not automatically be deleted.

You need to implement the changes to the Azure Cosmos DB for NoSQL account.

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

NOTE: Each correct selection is worth one point.

**Answer Area****Requirement**

Reads must return the most recent committed version of an item.

Items should not automatically be deleted.

**Change**

Change the consistency level by using the SDK.  
Change the default consistency level on the account.  
Configure the integrated cache on the account.  
Configure the account to write data to multiple Azure regions.

Change the Time to Live (TTL) property on items to 0.  
Change the Time to Live (TTL) property on items to -1.  
Configure a resource lock on the account.  
Configure a resource lock on the container.

**Answer Area****Requirement**

Reads must return the most recent committed version of an item.

Items should not automatically be deleted.

**Change**

Change the consistency level by using the SDK.  
Change the default consistency level on the account.  
Configure the integrated cache on the account.  
Configure the account to write data to multiple Azure regions.

Change the Time to Live (TTL) property on items to 0.  
Change the Time to Live (TTL) property on items to -1.  
Configure a resource lock on the account.  
Configure a resource lock on the container.

**Suggested Answer:**

**Jay456** Highly Voted 7 months, 1 week ago

\* The most recent committed version from any region - strong consistency. This cannot be changed by de SDK so should be done on the account level.

\* TTL = -1

upvoted 8 times

**brexamttopics** Most Recent 7 months, 1 week ago

Reads must return the most recent committed version: Change the default consistency level on the account; because, Consistency can only be relaxed at the SDK instance or request level. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account. And, Strong consistency offers a linearizability guarantee. Linearizability refers to serving requests concurrently. The reads are guaranteed to return the most recent committed version of an item.

upvoted 2 times



  **alvingeo** 8 months ago

For the requirement to always return the most recent committed version of an item, you should change the consistency level using the SDK to ensure strong consistency as needed. For preventing automatic deletion of items, set the Time to Live (TTL) property to 0. This will keep the items from expiring automatically.

So answer will be: 1 box. change the consistency level using the SDK\*

2 box : set the Time to Live (TTL) property to 0

upvoted 2 times


  **examtopicsLogin123** 6 months, 2 weeks ago

You can't set stronger consistency level via SDK.

"Consistency can only be relaxed at the SDK instance or request level. To move from weaker to stronger consistency, update the default consistency for the Azure Cosmos DB account."

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#override-the-default-consistency-level>

upvoted 1 times

  **NetoZee** 8 months ago

2 box is Time to live -1, because when it's -1 never expired.

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/time-to-live>

upvoted 2 times

## HOTSPOT

-


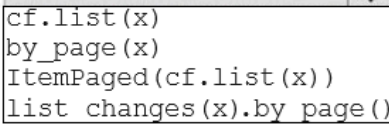
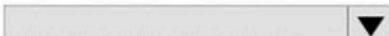
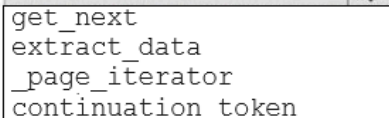
You are developing an application that monitors data added to an Azure Blob storage account.

You need to process each change made to the storage account.


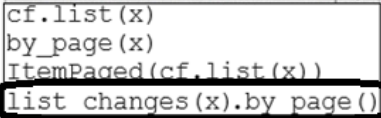

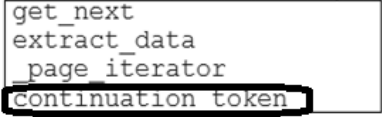
How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```
cf = ChangeFeedClient("", "")
x = None
while True:
    change_feed = cf. 
    
    for c in change_feed:
        ProcessChanges(c)
    x = change_feed. 
    
```

**Answer Area****Suggested Answer:**

```
cf = ChangeFeedClient("", "")
x = None
while True:
    change_feed = cf. 
    
    for c in change_feed:
        ProcessChanges(c)
    x = change_feed. 
    
```

 **overhill** 8 months ago

Correct, the image lacks underscore.

samples with continuation token here:

upvoted 2 times

 **alvingeo** 8 months ago

Given answer is correct! To process each change in the Azure Blob storage account, use `cf.list_changes(x).by_page()` to retrieve changes in pages. Then, set `x = change_feed.continuation_token` to keep track of the last processed position.

upvoted 4 times

You manage an Azure Cosmos DB for a NoSQL API account named account1. The account contains a database named db1, which contains a container named container1. You configure account1 with a session consistency level.

You plan to develop an application named App1 that will access container1. Individual instances of App1 must perform reads and writes. App1 must allow multiple nodes to participate in the same session.

You need to configure an object to share the session token between the nodes.


Which object should you use?

- A. Document response
- B. Request options -
- C. Feed options
- D. Connection policy

**Suggested Answer: A**

*Community vote distribution*

A (100%)

 **czegi90** 6 months, 1 week ago

**Selected Answer: B**

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#utilize-session-tokens>

upvoted 2 times

 **examtopicsLogin123** 6 months, 2 weeks ago

**Selected Answer: B**

In .NET SDK V2 it is: RequestOptions.SessionToken

in .NET SDK V3 it is: ItemRequestOptions.SessionToken

"If you wanted these nodes to participate in the same session ... you would have to send the SessionToken from FeedResponse<T> of the write action to the end-user ... to the CosmosClient for subsequent reads."

"To manage session tokens manually, get the session token from the response and set them per request."

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency#utilize-session-tokens>

FeedOptions also have the SessionToken property

<https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.feedoptions.sessiontoken>

upvoted 3 times

 **Mattt** 7 months, 3 weeks ago

**Selected Answer: B**

B. Request option

upvoted 1 times

 **Vichu\_1607** 7 months, 3 weeks ago

**Selected Answer: B**



B. Request options: The RequestOptions object allows you to specify options for individual requests, including the session token. By setting the session token in the RequestOptions, you can ensure that multiple nodes in your application share the same session token, thereby maintaining session consistency across those nodes.

upvoted 1 times

 **overhill** 8 months ago

For me it is Request Options the object we have to set with the session token

example here: <https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.requestoptions.sessiontoken?view=azure-dotnet>  
upvoted 3 times

  **passme847** 8 months, 2 weeks ago

**Selected Answer: A**

<https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.requestoptions.sessiontoken?view=azure-dotnet>  
upvoted 2 times

## DRAG DROP

-

## Case study

-

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

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

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

## To start the case study

-

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

## Background

-

Fourth Coffee is a global coffeehouse chain and coffee company recognized as one of the world's most influential coffee brands. The company is renowned for its specialty coffee beverages, including a wide range of espresso-based drinks, teas, and other beverages. Fourth Coffee operates thousands of stores worldwide.

## Current environment

-

The company is developing cloud-native applications hosted in Azure.

## Corporate website

-

The company hosts a public website located at <http://www.fourthcoffee.com/>. The website is used to place orders as well as view and update inventory items.

## Inventory items

-

In addition to its core coffee offerings, Fourth Coffee recently expanded its menu to include inventory items such as lunch items, snacks, and merchandise. Corporate team members constantly update inventory. Users can customize items. Corporate team members configure inventory items and associated images on the website.

## Orders

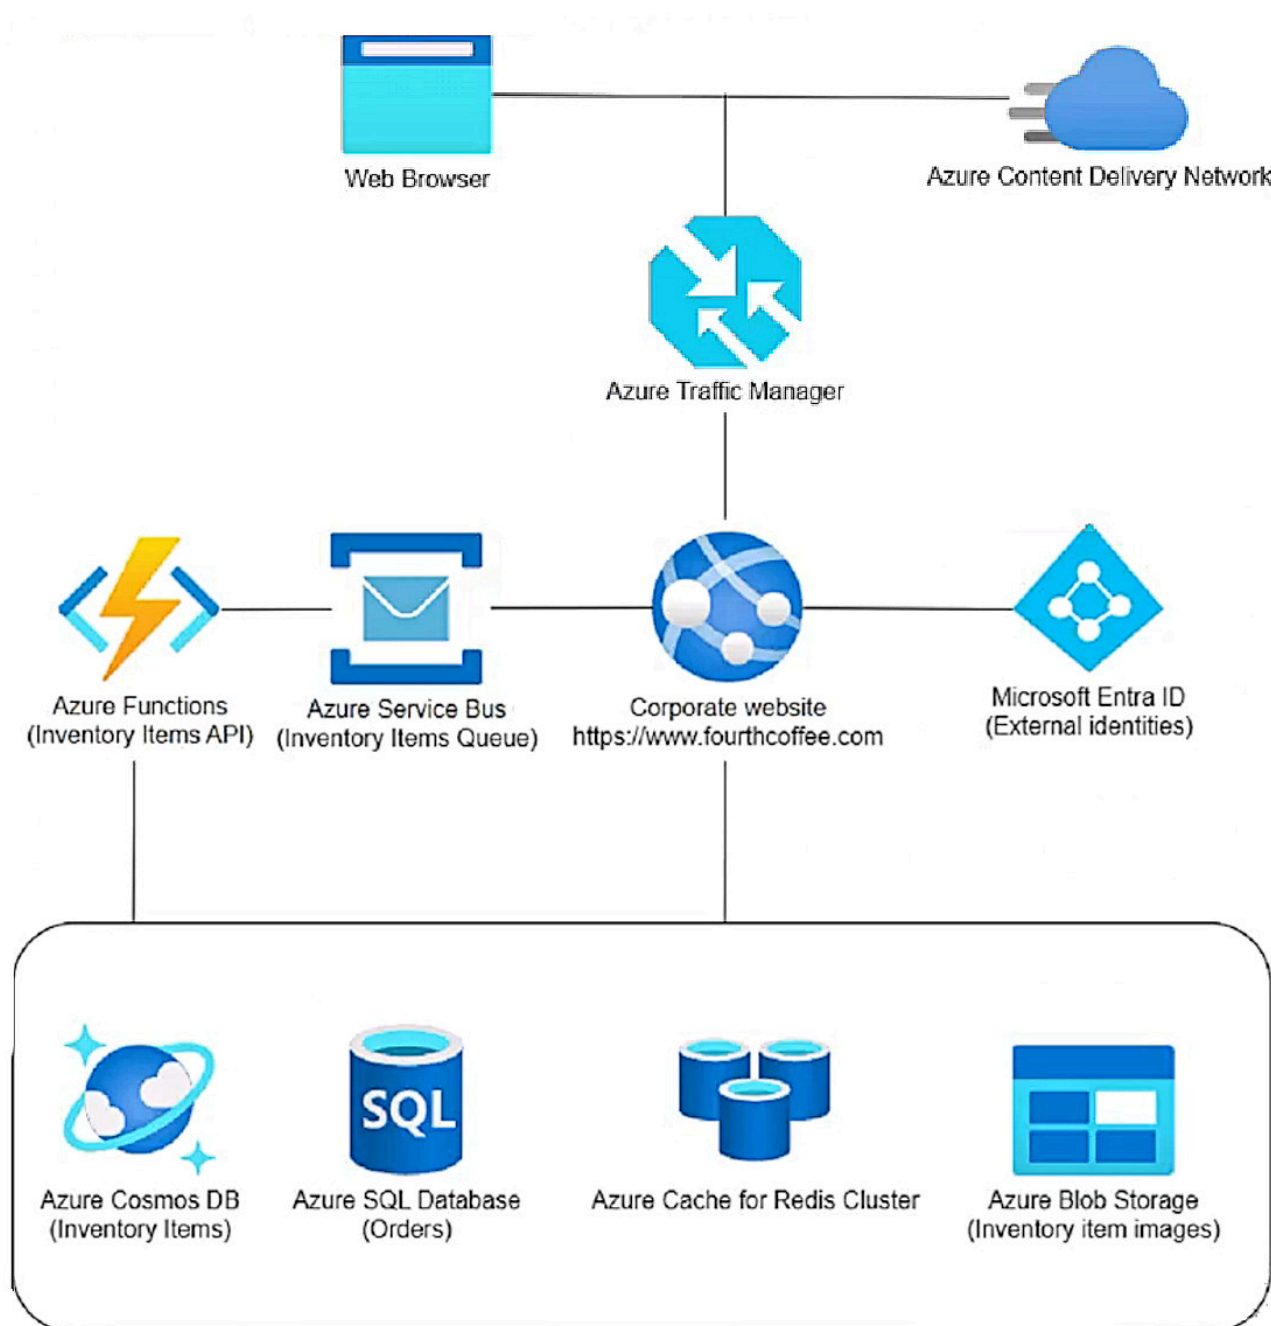
-

Associates in the store serve customized beverages and items to customers. Orders are placed on the website for pickup.

The application components process data as follows:

1. Azure Traffic Manager routes a user order request to the corporate website hosted in Azure App Service.
2. Azure Content Delivery Network serves static images and content to the user.
3. The user signs in to the application through a Microsoft Entra ID for customers tenant.
4. Users search for items and place an order on the website as item images are pulled from Azure Blob Storage.
5. Item customizations are placed in an Azure Service Bus queue message.
6. Azure Functions processes item customizations and saves the customized items to Azure Cosmos DB.
7. The website saves order details to Azure SQL Database.
8. SQL Database query results are cached in Azure Cache for Redis to improve performance.

The application consists of the following Azure services:



Requirements

The application components must meet the following requirements:

- Azure Cosmos DB development must use a native API that receives the latest updates and stores data in a document format.
- Costs must be minimized for all Azure services.
- Developers must test Azure Blob Storage integrations locally before deployment to Azure. Testing must support the latest versions of the Azure Storage APIs.

#### Corporate website

-

- User authentication and authorization must allow one-time passcode sign-in methods and social identity providers (Google or Facebook).
- Static web content must be stored closest to end users to reduce network latency.

#### Inventory items

-

- Customized items read from Azure Cosmos DB must maximize throughput while ensuring data is accurate for the current user on the website.
- Processing of inventory item updates must automatically scale and enable updates across an entire Azure Cosmos DB container.
- Inventory items must be processed in the order they were placed in the queue.
- Inventory item images must be stored as JPEG files in their native format to include exchangeable image file format (data) stored with the blob data upon upload of the image file.
- The Inventory Items API must securely access the Azure Cosmos DB data.

#### Orders

-

- Orders must receive inventory item changes automatically after inventory items are updated or saved.

#### Issues

-

- Developers are storing the Azure Cosmos DB credentials in an insecure clear text manner within the Inventory Items API code.
- Production Azure Cache for Redis maintenance has negatively affected application performance.

You need to store inventory item images.

Which Azure Blob Storage feature should you use? To answer, move the appropriate Azure Blob Storage features to the correct requirements. You may use each Azure Blob Storage feature once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

#### Answer Area

##### Azure Blob Storage features

Index tags
System properties
Change feed
User-defined metadata

##### Azure Blob Storage requirements

###### Requirement

Store images in native format

Store image Exif data

###### Azure Blob Storage feature



Suggested Answer:	Azure Blob Storage requirements	
	Requirement	Azure Blob Storage feature
	Store images in native format	Index tags
	Store image Exif data	User-defined metadata

  **tracknamepending** Highly Voted 8 months, 1 week ago

For the requirement "store images in native format", wouldn't "system properties" be more appropriate? None of these features affect the format the data is stored in, however setting the ContentType property to "image/jpeg" would indicate the native format of the stored binary data.

"User-defined metadata" seems correct for storing EXIF metadata.

upvoted 9 times

  **overhill** 8 months ago

I agree with this. System properties seems more appropriate.



I don't understand 100% what native format is but I'm supposing we are storing the files in the original format

upvoted 2 times

  **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

  **Curl8012** 5 months, 1 week ago

Should be

- System properties

- User-defined data

Link: <https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.documents.client.requestoptions.sessiontoken?view=azure-dotnet>

upvoted 3 times

## DRAG DROP

-

You have a static website hosted in an Azure Storage account named storage1. You access the website by using a URL that ends with the web.core.windows.net suffix.

You plan to configure the website to be accessible through the URL www.contoso.com. The website must be accessible during configuration.

The contoso.com zone is hosted in Azure DNS.

You need to complete the website configuration.

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

## Domain name configuration steps

Create a CNAME record as verify.www.contoso.com in Azure DNS.

Set the custom domain name of storage1 to www.contoso.com.

Identify the host name of the storage endpoint of the website.

Create an A record as verify.www.contoso.com in Azure DNS.

Create a CNAME record www.contoso.com in Azure DNS.

## Answer Area

1

2

3

4

## Suggested Answer:

## Domain name configuration steps

Create a CNAME record as verify.~~www~~.contoso.com in Azure DNS.

Set the custom domain name of ~~storage1~~ to www.contoso.com.

Identify the host name of the storage endpoint of the website.

Create an A record as verify.www.contoso.com in Azure DNS.

Create a CNAME record ~~www~~.contoso.com in Azure DNS.


## Answer Area

1

2

3

4

 **21c5438** 5 months, 2 weeks ago

Provided answer is correct.

The description says "The website must be accessible during configuration", because of this we have to "Map a custom domain with zero downtime".

Reference:

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-custom-domain-name?tabs=azure-portal#map-a-domain>

upvoted 4 times

 **c75e123** 1 month ago

I agree. This is easier to read:

1. Identify the host name of the storage endpoint of the website.
2. Create a CNAME record as verify.www.contoso.com in Azure DNS.
3. Set the custom domain name of storage1 to www.contoso.com
4. Create a CNAME record www.contoso.com in Azure DNS

upvoted 1 times

 **21c5438** 5 months, 2 weeks ago

Reference correction:

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-custom-domain-name?tabs=azure-portal#map-a-custom-domain-with-zero-downtime>

upvoted 1 times

You have an Azure Queue Storage account that contains a queue named queue1.

You plan to use Azure SDK for .NET to develop a solution that uses queue1.

You need to author C# code that will return an approximate number of messages in queue1. Your solution must minimize the development effort.

Which method should you use in your code?

- A. GetProperties method of the QueueClient class
- B. GetProperties method of the QueueServiceClient class
- C. PeekMessages method of the QueueClient class
- D. GetStatistics method of the QueueServiceClient class

**Suggested Answer: A**

*Community vote distribution*

A (100%)

 **GrzegorzZima** Highly Voted 6 months ago

**Selected Answer: A**

Looks correct: <https://learn.microsoft.com/en-us/azure/storage/queues/storage-quickstart-queues-dotnet?tabs=passwordless%2Croles-azure-portal%2Cenvironment-variable-windows%2Csign-in-azure-cli#get-the-queue-length>

upvoted 6 times

 **wafa\_chaari** Most Recent 5 months ago

**Selected Answer: A**

correct

upvoted 3 times

You manage an Azure Storage account named storage1.

You plan to load 1 million blobs into storage1.

You must assign key-value pairs to blobs so that both keys and their values are automatically indexed and searchable by using the built-in services of storage1.

You need to run the command to assign key-value pairs.

Which command should you run?

- A. Update -AzStorageBlobServiceProperty
- B. Set-AzStorageBlobTag
- C. az storage blob service-properties update
- D. Set-AzStorageBlobContent

**Suggested Answer: B**

*Community vote distribution*

B (100%)

🗳️ 👤 **J\_obando13** 1 month, 2 weeks ago

**Selected Answer: B**

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)  
upvoted 2 times

🗳️ 👤 **c01efe8** 6 months ago

**Selected Answer: B**

<https://stackoverflow.com/questions/70097392/what-is-the-optimal-way-of-setting-blob-index-tag-in-azure-blob-storage-for-larg>  
upvoted 4 times

🗳️ 👤 **c75e123** 2 months ago

I agree.

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-index-how-to?tabs=azure-powershell#get-set-and-update-blob-index-tags>  
upvoted 1 times

## DRAG DROP

-

You manage an Azure Cosmos DB for a NoSQL API account named account1. You configure account1 with the default consistency level.

An application named app1 must access containers in account1 to perform read and write operations. The connections from app1 to account1 must be established by using the direct mode.

You plan to configure app1 to override the default consistency level by using the Azure Cosmos DB SDK client.

You need to set the maximum consistency level for app1 to use for read and write operations.

Which consistency level should you set? To answer, move the appropriate maximum consistency levels to the correct operation types. You may use each maximum consistency level once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

## Consistency levels for operations

Strong
Session
Bounded staleness

## Consistency levels for operations

Operation type	Maximum consistency level
Reads	
Writes	

## Consistency levels for operations

## Suggested Answer:

Operation type	Maximum consistency level
Reads	Bounded staleness
Writes	Strong

**21c5438** Highly Voted 5 months, 1 week ago

The default consistency level is session.

Only reads consistency level can be overridden, but it can not be stronger than the consistency level set for the account.

Therefore:

session

session

upvoted 10 times

**c75e123** 2 months ago

Agree, here is the source:

<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-manage-consistency?tabs=portal%2Cdotnetv2%2Capi-async#configure-the-default-consistency-level>

upvoted 1 times

**florianwicher** 5 months ago

Agree, for the exact same reason

upvoted 1 times

**vizay** Most Recent 1 month, 4 weeks ago

Writes are always done with full consistency within the write region(s), regardless of the consistency level you choose.

so the given answer is correct

upvoted 1 times

**vizay** 1 month, 4 weeks ago

By setting the consistency level in app1 to Session, you ensure that the application operates at a consistency level that is supported by the underlying account configuration. If you tried to set it to a stronger consistency level like Strong, the Azure Cosmos DB service would enforce the account's default consistency (Session).

session and session

upvoted 1 times

  **vizay** 1 month, 4 weeks ago

Account-level consistency is set to Session.

app1 tries to override the consistency level using the SDK.

For reads: app1 can set the consistency to Session, Eventual, or Consistent Prefix, but not Strong.

For writes: app1 cannot set a stronger consistency than Session either, because the account default restricts it to Session for both reads and writes.

upvoted 1 times

## DRAG DROP

You manage an Azure subscription associated with a Microsoft Entra tenant named contoso.com. The subscription contains an Azure Blob Storage account named storage1. Your user account has the Contributor Azure role-based access control (RBAC) role within the scope of the subscription.

You plan to implement secure access to containers and blobs in storage1. Your solution must satisfy the following requirements:

- Authorization requests to access storage1 content must be authenticated by using Microsoft Entra credentials.
- Authorized access to storage1 content must be time-limited based on arbitrary values specified when requests are raised.
- The principle of least privilege must be satisfied.

You need to implement the plan.

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

**Secure access implementation**

Assign the Owner role to your user account at the scope of the storage account.

Request a user delegation key.

Generate a shared access signature token.

Configure a stored access policy.

Acquire an OAuth 2.0 token from Microsoft Entra ID.

**Answer Area****Answer Area**

Acquire an OAuth 2.0 token from Microsoft Entra ID.

**Suggested Answer:**

Request a user delegation key.

Generate a shared access signature token.

**vizay** 1 month, 4 weeks ago

Assign the Owner role to your user account at the scope of the storage account. (While "Storage Blob Data Contributor" is the least privilege for accessing data, the ability to generate a User Delegation SAS might require broader permissions, potentially including those within "Owner" or a role that allows acting as the data owner for authorization purposes. Let's proceed with this option as it's available, and we can discuss potential least privilege refinements afterward if needed for a different scenario.)

Request a user delegation key. (This is a necessary step to create a User Delegation SAS using your Azure AD credentials.)

Generate a shared access signature token. (This will be the User Delegation SAS, leveraging the key obtained in the previous step and the permissions granted by your RBAC role.)



upvoted 1 times

**21c5438** 5 months, 1 week ago

1. Assign the Owner role
2. Acquire an OAuth 2.0 token
3. Request a user delegation key



upvoted 2 times

  **21c5438** 5 months, 1 week ago

Correction: provided answer is correct.

The task says that we already have the Contributor role.

upvoted 5 times

### Case study -

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

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

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

### To start the case study -

To display the first question in this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. When you are ready to answer a question, click the Question button to return to the question.

### Background -

Fourth Coffee is a global coffeehouse chain and coffee company recognized as one of the world's most influential coffee brands. The company is renowned for its specialty coffee beverages, including a wide range of espresso-based drinks, teas, and other beverages. Fourth Coffee operates thousands of stores worldwide.

### Current environment -

The company is developing cloud-native applications hosted in Azure.

### Corporate website -

The company hosts a public website located at <http://www.fourthcoffee.com/>. The website is used to place orders as well as view and update inventory items.

### Inventory items -

In addition to its core coffee offerings, Fourth Coffee recently expanded its menu to include inventory items such as lunch items, snacks, and merchandise. Corporate team members constantly update inventory. Users can customize items. Corporate team members configure inventory items and associated images on the website.

### Orders -

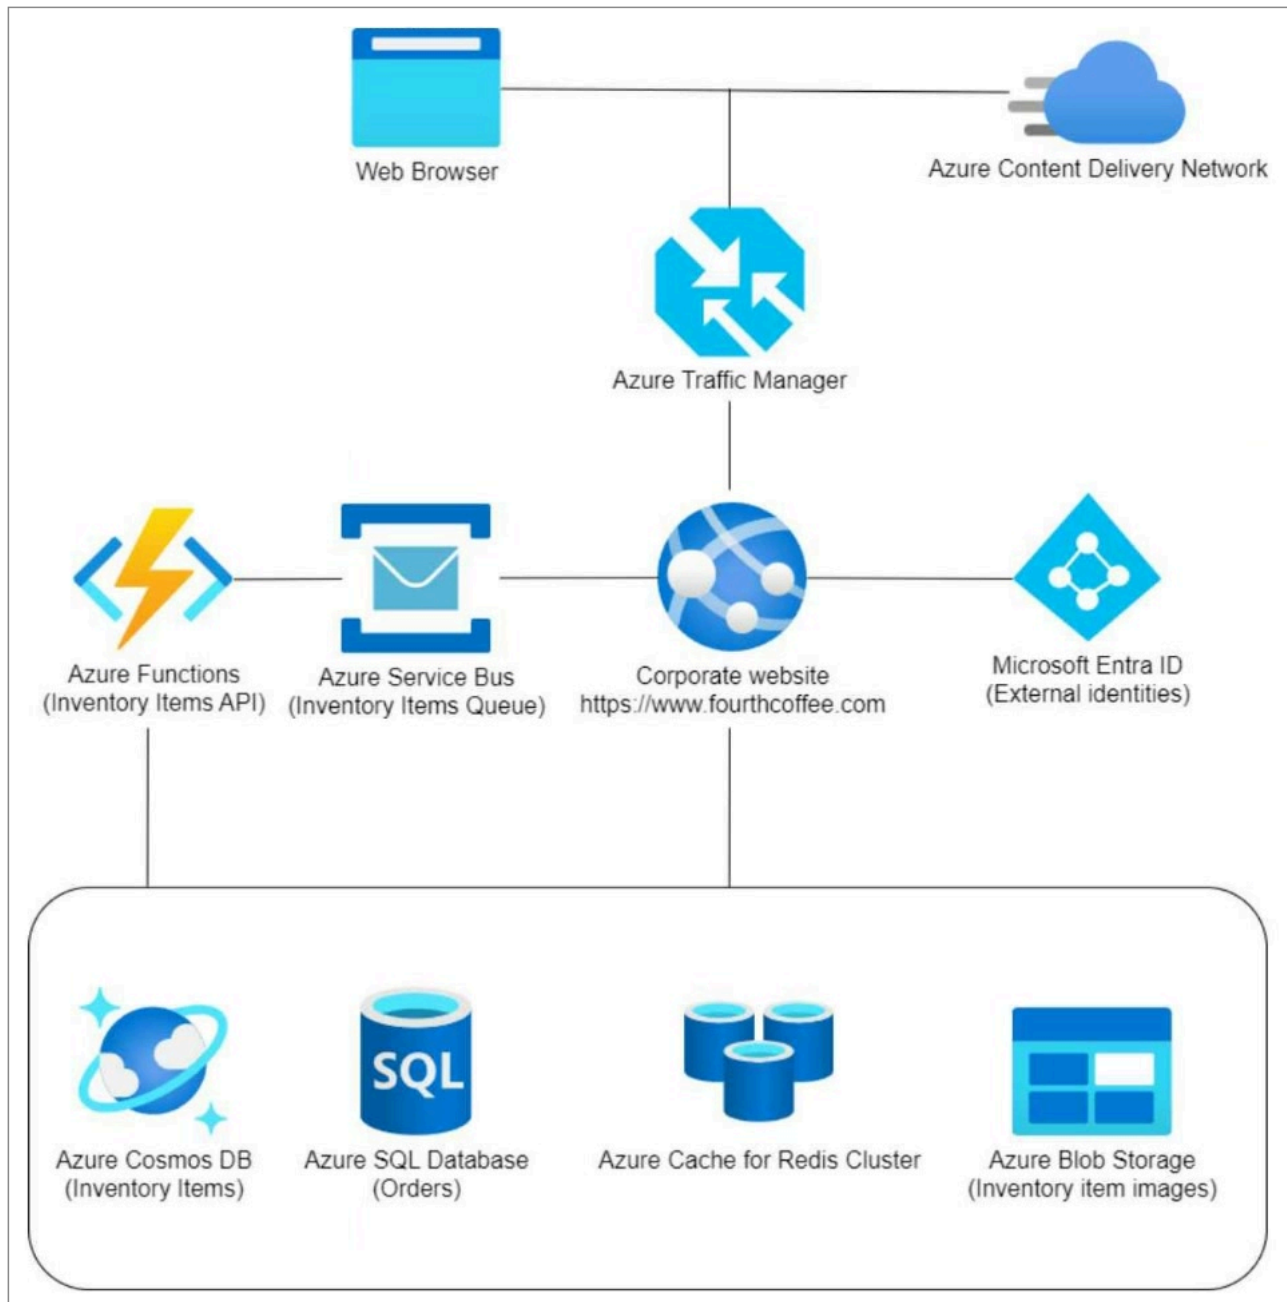
Associates in the store serve customized beverages and items to customers. Orders are placed on the website for pickup.

The application components process data as follows:

1. Azure Traffic Manager routes a user order request to the corporate website hosted in Azure App Service.
2. Azure Content Delivery Network serves static images and content to the user.
3. The user signs in to the application through a Microsoft Entra ID for customers tenant.
4. Users search for items and place an order on the website as item images are pulled from Azure Blob Storage.
5. Item customizations are placed in an Azure Service Bus queue message.
6. Azure Functions processes item customizations and saves the customized items to Azure Cosmos DB.
7. The website saves order details to Azure SQL Database.

8. SQL Database query results are cached in Azure Cache for Redis to improve performance.

The application consists of the following Azure services:



#### Requirements -

The application components must meet the following requirements:

- Azure Cosmos DB development must use a native API that receives the latest updates and stores data in a document format.
- Costs must be minimized for all Azure services.
- Developers must test Azure Blob Storage integrations locally before deployment to Azure. Testing must support the latest versions of the Azure Storage APIs.

#### Corporate website -

- User authentication and authorization must allow one-time passcode sign-in methods and social identity providers (Google or Facebook).
- Static web content must be stored closest to end users to reduce network latency.

#### Inventory items -

- Customized items read from Azure Cosmos DB must maximize throughput while ensuring data is accurate for the current user on the website.
- Processing of inventory item updates must automatically scale and enable updates across an entire Azure Cosmos DB container.
- Inventory items must be processed in the order they were placed in the queue.
- Inventory item images must be stored as JPEG files in their native format to include exchangeable image file format (data) stored with the blob data upon upload of the image file.
- The Inventory Items API must securely access the Azure Cosmos DB data.

#### Orders -

- Orders must receive inventory item changes automatically after inventory items are updated or saved.

#### Issues -

- Developers are storing the Azure Cosmos DB credentials in an insecure clear text manner within the Inventory Items API code.
- Production Azure Cache for Redis maintenance has negatively affected application performance.

You need to support local development testing for developers.

Which tool should you use?

- A. Azurite
- B. Azure Storage Emulator
- C. SQL Server Management Studio (SSMS)
- D. Azure Storage Explorer

**Suggested Answer: A**

*Community vote distribution*



**21c5438** Highly Voted 5 months, 1 week ago

**Selected Answer: A**

Correct answer is Azurite, because Azure Storage Emulator is deprecated:

<https://learn.microsoft.com/en-us/azure/storage/common/storage-use-azurite?tabs=visual-studio%2Cblob-storage>

<https://learn.microsoft.com/en-us/azure/storage/common/storage-use-emulator>

upvoted 6 times

HOTSPOT

-

You have an Azure Blob Storage account named account1.

You plan to grant permissions for access to blobs in account1 by using a combination of a shared access signature token and a stored access policy. You must set the duration of the token validity by using the stored access policy.

You need to complete the configuration of the stored access policy and generate the shared access signature token.

Which token setting and token type should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

### Answer Area

#### Storage access configuration

Configuration setting	Value
Token setting to configure	<div>Protocol Source IP Permissions</div>
Token type to generate	<div>Account shared access signatures Blob container shared access signatures User delegation shared access signatures</div>

Suggested Answer:

### Answer Area

#### Storage access configuration

Configuration setting	Value
Token setting to configure	<div>Protocol Source IP <b>Permissions</b></div>
Token type to generate	<div>Account shared access signatures <b>Blob container shared access signatures</b> User delegation shared access signatures</div>

21c5438 5 months, 1 week ago

Provided answer is correct.

Access policy can only be set on a service SAS.

Regarding the configuration:

<https://learn.microsoft.com/en-us/azure/ai-services/translator/document-translation/how-to-guides/create-sas-tokens?tabs=Containers#create-sas-tokens-in-the-azure-portal>

upvoted 2 times

You are developing a Java application that uses Cassandra to store key and value data. You plan to use a new Azure Cosmos DB resource and the Cassandra API in the application. You create an Azure Active Directory (Azure AD) group named Cosmos DB Creators to enable provisioning of Azure Cosmos accounts, databases, and containers.

The Azure AD group must not be able to access the keys that are required to access the data.

You need to restrict access to the Azure AD group.

Which role-based access control should you use?

- A. DocumentDB Accounts Contributor
- B. Cosmos Backup Operator
- C. Cosmos DB Operator
- D. Cosmos DB Account Reader

**Suggested Answer: C**

Azure Cosmos DB now provides a new RBAC role, Cosmos DB Operator. This new role lets you provision Azure Cosmos accounts, databases, and containers, but can't access the keys that are required to access the data. This role is intended for use in scenarios where the ability to grant access to Azure Active Directory service principals to manage deployment operations for Cosmos DB is needed, including the account, database, and containers.

Reference:

<https://azure.microsoft.com/en-us/updates/azure-cosmos-db-operator-role-for-role-based-access-control-rbac-is-now-available/>

Community vote distribution

C (100%)

 **mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: C

DocumentDB Account Contributor: Can manage Azure Cosmos DB accounts.

Cosmos Backup Operator: Can submit a restore request for Azure portal for a periodic backup enabled database or a container. Can modify the backup interval and retention on the Azure portal. Cannot access any data or use Data Explorer.

Cosmos DB Operator: Can provision Azure Cosmos accounts, databases, and containers. Cannot access any data or use Data Explorer.

Cosmos DB Account Reader: Can read Azure Cosmos DB account data.

CosmosRestoreOperator: Can perform restore action for Azure Cosmos DB account with continuous backup mode.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/role-based-access-control>

upvoted 61 times

 **OPT\_001122** 2 years, 10 months ago

very good and helpful explanation

upvoted 2 times

 **SachinV** 1 year, 7 months ago


looks correct, in Nov182023, scored 962

upvoted 1 times

 **RaviKS** Highly Voted 4 years, 6 months ago

Answer is correct

upvoted 38 times

 **J\_obando13** Most Recent 1 month, 2 weeks ago

**Selected Answer: C**

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

🗒️ 👤 **overhill** 8 months ago

Roles

<https://learn.microsoft.com/en-us/azure/role-based-access-control/built-in-roles?context=%2Fazure%2Fcosmos-db%2Fnosql%2Fsecurity%2Fcontext%2Fcontext#databases>

upvoted 1 times

🗒️ 👤 **Vichu\_1607** 8 months, 2 weeks ago

**Selected Answer: C**

Cosmos DB Operator

upvoted 1 times

🗒️ 👤 **Vichu\_1607** 8 months, 3 weeks ago

**Selected Answer: A**

DocumentDB Account Contributor:

upvoted 1 times

🗒️ 👤 **Vichu\_1607** 8 months, 3 weeks ago

Answer is C

upvoted 1 times

🗒️ 👤 **azpardha** 1 year, 2 months ago

Got this on 04/07/2024, Went with highly voted answer, scored:820

upvoted 2 times

🗒️ 👤 **Woksi** 1 year, 4 months ago

Got on 07/02/2024

upvoted 1 times

🗒️ 👤 **richardc7** 1 year, 9 months ago

got 28/09/2023

upvoted 1 times

🗒️ 👤 **richardc7** 1 year, 9 months ago

got on 28/09/2023

upvoted 1 times

🗒️ 👤 **mihailos** 1 year, 10 months ago

Got it in exam 28/08/23. Went with proposed answer. Scored 912

upvoted 2 times

🗒️ 👤 **basquiat** 1 year, 10 months ago

Got this one in 2023-08-08

upvoted 3 times

🗒️ 👤 **JH81** 2 years ago

**Selected Answer: C**

Got this on 6/28/2023 and passed with 850. Went with answer.

upvoted 3 times

🗒️ 👤 **aragones** 2 years, 1 month ago

Got this 2023-05-12.

My case:

case: You need to configure authorization.

case: You need to ensure the app does not time out and processes the blob data



case study: VanArsdel Inc Canada

upvoted 2 times

🗒️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 1 times

  **Esward** 2 years, 5 months ago

Given answer is correct

<https://azure.microsoft.com/en-us/updates/azure-cosmos-db-operator-role-for-role-based-access-control-rbac-is-now-available/>  
upvoted 2 times



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 are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.

You need to configure authorization.

Solution: Configure the Azure Web App for the website to allow only authenticated requests and require Azure AD log on.

Does the solution meet the goal?

A. Yes

B. No


**Suggested Answer: B**

Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

Community vote distribution



**fadikh** Highly Voted 4 years, 3 months ago

Answer is correct. The proposed solution lacks the authorization part.  
upvoted 33 times

**ferut** 4 years, 1 month ago

Aggree. Setting up the access on AD is one thing. The application should be set up to differentiate different roles and access.  
upvoted 5 times

**edengoforit** 3 years, 5 months ago

Answer is NO  
upvoted 3 times

**kondapaturi** Highly Voted 4 years ago

Answer – No, Here you need to create an application in Azure AD. Then set the groupMembershipClaims claims. Then inspect the token in the application to see if the user is part of that group.  
Hence - B is correct  
upvoted 14 times

**Vichu\_1607** Most Recent 8 months, 3 weeks ago

Selected Answer: B  
No, the solution does not meet the goal.

While configuring the Azure Web App to allow only authenticated requests and require Azure AD log on is a necessary step for securing the application, it does not address the requirement of assigning users one of the specific permission levels (admin, normal, reader) based on their Azure AD group membership.

To meet this requirement, you would need to implement role-based access control (RBAC) in your application. This would involve mapping Azure AD groups to roles within your application, and then assigning permissions to those roles.  
upvoted 2 times

**Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850  
upvoted 3 times

**Tarajee** 1 year, 9 months ago

On my exam 2023sept

upvoted 1 times

🗨️ 👤 **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 4 times

🗨️ 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Chosen No. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 3 times

🗨️ 👤 **Esward** 2 years, 5 months ago

Given answer is correct

<https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups>

upvoted 1 times

🗨️ 👤 **capable** 2 years, 7 months ago

Got this case study in my exam 27/11/2022

upvoted 2 times

🗨️ 👤 **gmishra88** 2 years, 8 months ago

If you think about it having the groups claim is also not enough. You need to check for it. So, even if the group claims was there it could be no.

Because of that reason it could have been yes, because you need to do that step as well as groups claim and also check if the claim is authorized.

These kind of questions are just Microsoft way of confusion

upvoted 2 times

🗨️ 👤 **bhakk** 2 years, 12 months ago

Answer is B

upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

**Selected Answer: B**

B is correct answer

upvoted 1 times

🗨️ 👤 **pradipbobbhate** 3 years, 3 months ago

**Selected Answer: A**

correct Answer

upvoted 1 times

🗨️ 👤 **Freidrich** 3 years, 4 months ago

**Selected Answer: B**

The correct answer is B: No.

upvoted 1 times

🗨️ 👤 **edengoforit** 3 years, 5 months ago

This is the correct answer:

Solution:

⇒ Create a new Azure AD application. In the application's manifest, set value of the groupMembershipClaims option to All.

⇒ In the website, use the value of the groups claim from the JWT for the user to determine permissions.

upvoted 5 times

🗨️ 👤 **edengoforit** 3 years, 5 months ago

Provided answer and explanation is correct.

upvoted 1 times

🗨️ 👤 **ehurfheiz** 3 years, 5 months ago

**Selected Answer: B**

B seems to be the correct answer

upvoted 1 times

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 are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.

You need to configure authorization.

Solution:

- ⇒ Create a new Azure AD application. In the application's manifest, set value of the groupMembershipClaims option to All.
- ⇒ In the website, use the value of the groups claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: A

To configure Manifest to include Group Claims in Auth Token

1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:
2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
3. Locate the `groupMembershipClaims` setting. Set its value to either `SecurityGroups` or `All`. To help you decide which:
  - ⇒ `SecurityGroups` - groups claim will contain the identifiers of all security groups of which the user is a member.
  - ⇒ `All` - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

Now your application will include group claims in your manifest and you can use this fact in your code.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

Community vote distribution

A (100%)

**mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: Yes

Reference:

<https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups>

upvoted 47 times

**kondapaturi** Highly Voted 4 years ago

Answer – Yes, this is the correct approach. The Application manifest has a setting for groupMembershipClaims. By setting this to all, the Azure AD groups that the user belongs to will be returned as part of the claims in the JWT token.

upvoted 9 times

**Vichu\_1607** Most Recent 8 months, 3 weeks ago

Selected Answer: A

Yes, the solution does not meet the goal.

While configuring the Azure Web App to allow only authenticated requests and require Azure AD log on is a necessary step for securing the application, it does not address the requirement of assigning users one of the specific permission levels (admin, normal, reader) based on their Azure AD group membership.

To meet this requirement, you would need to implement role-based access control (RBAC) in your application. This would involve mapping Azure AD groups to roles within your application, and then assigning permissions to those roles.

upvoted 1 times

**Dixavado** 1 year, 9 months ago

Selected Answer: A

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 1 times

🗳️ 👤 **Tarajee** 1 year, 9 months ago

On my exam 2023sept

upvoted 1 times

🗳️ 👤 **NightshadeRC** 1 year, 11 months ago

Had this question in today's exam: 2023-07-26

upvoted 2 times

🗳️ 👤 **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 2 times

🗳️ 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Chosen yes. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

🗳️ 👤 **BrettusMaximus** 2 years, 4 months ago

B

Sure, A would technically work but not the best solution as unauthorized users can hit the website.

Best practice is to use groupMembershipClaims in the App Registration to restrict access at the App level.

upvoted 2 times

🗳️ 👤 **Esward** 2 years, 5 months ago

Given answer A is correct as per MS docs

<https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups>

upvoted 2 times

🗳️ 👤 **serpevi** 2 years, 9 months ago

Got this in 09/22 , went just with this as YES, score 927.

upvoted 4 times

🗳️ 👤 **Eltooth** 3 years ago

**Selected Answer: A**

A is correct answer.

upvoted 3 times

🗳️ 👤 **PieroFranco** 3 years, 3 months ago

**Selected Answer: A**

The answer is correct. If you do not add the groupMembershipClaims to the manifest you wont see any AD group in the auth token.

upvoted 2 times

🗳️ 👤 **Freidrich** 3 years, 4 months ago

**Selected Answer: A**

The answer is correct.

upvoted 1 times

🗳️ 👤 **sujitwarrier11** 4 years ago

I think the answer is no. group claims are at Azure Ad tenant level. The question mentions that the authorization should be for the app. So roles would be the better option here.

upvoted 1 times

🗳️ 👤 **ZodiaC** 3 years, 11 months ago

thats not true, look link plz: <https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups>

upvoted 1 times

🗳️ 👤 **UnknowMan** 4 years, 1 month ago

Answer is correct

upvoted 1 times

🗳️ 👤 **glam** 4 years, 1 month ago

A. Yes

upvoted 1 times



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 are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.

You need to configure authorization.

Solution:

- ⇒ Create a new Azure AD application. In the application's manifest, define application roles that match the required permission levels for the application.
- ⇒ Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: B

To configure Manifest to include Group Claims in Auth Token


1. Go to Azure Active Directory to configure the Manifest. Click on Azure Active Directory, and go to App registrations to find your application:
2. Click on your application (or search for it if you have a lot of apps) and edit the Manifest by clicking on it.
3. Locate the `groupMembershipClaims` setting. Set its value to either `SecurityGroups` or `All`. To help you decide which:
  - ⇒ `SecurityGroups` - groups claim will contain the identifiers of all security groups of which the user is a member.
  - ⇒ `All` - groups claim will contain the identifiers of all security groups and all distribution lists of which the user is a member

Now your application will include group claims in your manifest and you can use this fact in your code.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

Community vote distribution



🗳️ 👤 [Removed] Highly Voted 👍 4 years, 5 months ago

I agree that this solution should work as well. The roles get assigned by AD groups, so the requirement "A user's Azure AD group membership must be used to determine the permission level" is met.

This solution should be answered with "yes".

This scenario has 2 solutions provided as the approach using the "groupMembershipClaims" is possible as well.

That's OK as it says "Some question sets might have more than one correct solution, while others might not have a correct solution."

upvoted 63 times

🗳️ 👤 Mal22002 3 years, 3 months ago

There are many solutions when you're writing the code but the question is whether provided solution works or not and it will not work. The second part of the solution failed the requirement "A user's Azure AD group membership must be used to determine the permission level." as it's saying "use the value of roles claim".

The answer is No.

upvoted 9 times

🗳️ 👤 Adiu 2 years, 4 months ago

But AD groups are used to assign correct roles to them. So the roles claim is based on user's AD groups membership.

upvoted 3 times

🗳️ 👤 Adrian1405 Highly Voted 👍 4 years, 6 months ago

This should be the correct answer: <https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps>

upvoted 20 times

🗨️ 👤 **iiiihhh** 4 years, 6 months ago

I think that without the "groupMembershipClaims" setting you will not be able to obtain group membership information and so the proposal doesn't meet the goal.

upvoted 8 times

🗨️ 👤 **bugimachi** 4 years, 6 months ago

I am not too sure about this. "groupMembershipClaims" will include AAD groups' IDs in the JWT "group" array, which is one way to go.

The other way would be to add a app role in the manifest (which is not necessary for the first approach!) and then add groups to this role. This will include the app roles (but not the group IDs) in the JWT's "roles" array (but not in the "groups" array!).

I liked this blog post to understand the difference: <https://joonasw.net/view/using-groups-vs-using-app-roles-in-azure-ad-apps>

upvoted 13 times

🗨️ 👤 **clarionprogrammer** 4 years, 2 months ago

Excellent post! A. Yes.

Final answer.

upvoted 6 times

🗨️ 👤 **FatSeal** 2 years, 10 months ago

A user's Azure AD group membership must be used to determine the permission level!!!!!!!!!!!!!!

upvoted 2 times

🗨️ 👤 **warchoon** 2 years, 4 months ago

You can use gropes for roles

<https://learn.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps#:~:text=In%20the%20Create%20app%20role%20pane%2C%20enter%20the%20settings%20for%20the%20role.>

upvoted 1 times

🗨️ 👤 **ALittleBunny** 4 years, 6 months ago

You're totally right. It's just managing permissions in two different ways, groupMembershipClaims and assigning the appropriate Azure AD group to each role, which are both correct.

upvoted 7 times

🗨️ 👤 **ahadjithoma** 4 years, 6 months ago

Requirement: A user's Azure AD group membership must be used to determine the permission level.

Does that solution meets the above requirement?

upvoted 3 times

🗨️ 👤 **bugimachi** 4 years, 6 months ago

I think so; you can add groups to app roles defined in manifest.

upvoted 4 times

🗨️ 👤 **rdemontis** 4 years, 3 months ago

I think you can't instead, because AppRoles are intended for RBAC access control and not AAD group membership as required in the question. See this for more info: <https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps>

upvoted 2 times

🗨️ 👤 **Vichu\_1607** Most Recent 8 months, 3 weeks ago

Selected Answer: A

Yes, the solution meets the goal.

Defining application roles in the Azure AD application's manifest that match the required permission levels (admin, normal, reader) is a correct step. These roles can then be assigned to users or groups in Azure AD.

Assigning the appropriate Azure AD groups to these application roles will ensure that a user's group membership determines their permission level in the application. When a user signs in, their role memberships will be included in their token, and your application can use this information to authorize their actions.

Using the value of the roles claim from the JWT to determine permissions in the website is the correct approach. This claim will contain all the roles that the user is a member of.

upvoted 1 times

🗨️ 👤 **Jass1nonly** 1 year, 4 months ago

This seems a valid approach from the links shared in the discussions.

My answer is Yes. The given solution meets the goal.

upvoted 1 times

🗨️ 👤 **FeriAZ** 1 year, 4 months ago

**Selected Answer: A**

This approach is a standard practice for implementing role-based access control (RBAC) in applications that use Azure AD for authentication. It leverages Azure AD's built-in capabilities for managing users, groups, and roles, providing a scalable and manageable method for handling authorization based on group membership.

upvoted 2 times

🗨️ 👤 **AlbertoBT** 1 year, 7 months ago

**Selected Answer: B**

"Azure AD group membership must be used"

<https://learn.microsoft.com/en-us/entra/identity-platform/howto-add-app-roles-in-apps#app-roles-vs-groups>

App roles -> Provided in the roles claim

Groups -> Provided in groups claim.

Simple as that

upvoted 2 times

🗨️ 👤 **Weam** 1 year, 7 months ago

I will go for (NO)

because while adding app roles, we don't have to apply any changes in the manifest file. The answer would be correct if it didn't state that we will apply changes in the manifest file. So the correct answer for this case study is to use group claims and apply change to manifest file to set group claims to (ALL) or (Security Group) or whatever option

upvoted 1 times

🗨️ 👤 **Dixavado** 1 year, 9 months ago

**Selected Answer: B**

It was on my exam today (2023-09-26) I went with the examtopics answer - score 850

upvoted 1 times

🗨️ 👤 **ayush66203** 1 year, 8 months ago

These examtopics questions were sufficient for you to score 850? or you prepared from other sources too?

upvoted 1 times

🗨️ 👤 **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 2 times

🗨️ 👤 **Zoriark** 2 years, 2 months ago

ChatGPT

A. Yes, the solution meets the goal.

The solution correctly suggests creating a new Azure AD application and defining application roles that match the required permission levels. By assigning the appropriate Azure AD groups to each role, the application can determine a user's permission level based on their group membership. The roles claim from the JWT token can then be used in the website to determine the user's permissions. This approach allows for granular access control and ensures that users are assigned the correct permission level based on their group membership.

upvoted 2 times

🗨️ 👤 **70PineApple** 2 years, 4 months ago

Got this in exam today..20/02/23

upvoted 1 times

🗨️ 👤 **Jeeaaak** 2 years, 4 months ago

does the answer is yes

upvoted 2 times

🗨️ 👤 **Neckbread** 1 year, 5 months ago

yes answer is the does

upvoted 1 times

🗨️ 👤 **Esward** 2 years, 5 months ago

yes is the correct answer



<https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps>

upvoted 1 times

🗨️ 👤 **andrvelich** 3 years, 2 months ago

**Selected Answer: A**

should work as well

upvoted 1 times

🗨️ 👤 **Azprep** 3 years, 2 months ago

Answer should be yes

upvoted 1 times

🗨️ 👤 **Mal22002** 3 years, 3 months ago

**Selected Answer: B**

The answer is No, the explanation is useless though.

Requirement: A user's Azure AD group membership must be used to determine the permission level.

The solution that failed the requirement: Assign the appropriate Azure AD group to each role. In the website, use the value of the roles claim from the JWT for the user to determine permissions.

Azure AD group membership can be accessed in groups claim not in roles claim.

Reference: <https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-add-app-roles-in-azure-ad-apps>

upvoted 4 times

🗨️ 👤 **mmdex** 2 years, 4 months ago

Read carefully: "A user's Azure AD group membership must be used to determine the permission level."

You are not trying to determine user's group membership. You are trying to determine their permission level. You don't need groups claim, roles is enough.

upvoted 1 times

🗨️ 👤 **Freidrich** 3 years, 4 months ago

**Selected Answer: A**

I think the correct answer is A: Yes.

upvoted 3 times

🗨️ 👤 **idroj** 3 years, 4 months ago

**Selected Answer: A**

The correct answer is yes

upvoted 3 times

## DRAG DROP -

You are developing an application to securely transfer data between on-premises file systems and Azure Blob storage. The application stores keys, secrets, and certificates in Azure Key Vault. The application uses the Azure Key Vault APIs.

The application must allow recovery of an accidental deletion of the key vault or key vault objects. Key vault objects must be retained for 90 days after deletion.

You need to protect the key vault and key vault objects.

Which Azure Key Vault feature should you use? To answer, drag the appropriate features to the correct actions. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Features	Answer Area
Access policy	
Purge protection	
Soft delete	
Shared access signature	

Action	Feature
Enable retention period and accidental deletion.	Feature
Enforce retention period and accidental deletion.	Feature

## Suggested Answer:

Features	Answer Area
Access policy	
Purge protection	
Soft delete	
Shared access signature	

Action	Feature
Enable retention period and accidental deletion.	Soft delete
Enforce retention period and accidental deletion.	Purge protection

Box 1: Soft delete -

When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.

Box 2: Purge protection -

Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.


When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview>

 **andsol** Highly Voted 4 years, 3 months ago

The answer is correct.  
upvoted 64 times

 **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: Soft delete

When soft-delete is enabled, resources marked as deleted resources are retained for a specified period (90 days by default). The service further provides a mechanism for recovering the deleted object, essentially undoing the deletion.

This can be achieved with the help of the soft-delete feature of the key vault.

Box 2: Purge protection

Purge protection is an optional Key Vault behavior and is not enabled by default. Purge protection can only be enabled once soft-delete is enabled.

When purge protection is on, a vault or an object in the deleted state cannot be purged until the retention period has passed. Soft-deleted vaults and objects can still be recovered, ensuring that the retention policy will be followed.

This can be achieved with the help of the purge protection feature of the key vault.

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview>

upvoted 47 times

  **minh1302** 2 years, 11 months ago

Thank mlantonis, learn from you a lot ( AZ-104, AZ-204)

upvoted 5 times

  **Vichu\_1607** Most Recent 8 months, 3 weeks ago

The feature you should use to enable retention period and accidental deletion is Soft Delete.

Soft Delete, when enabled, allows recovery of the deleted key vault or key vault objects within a specified retention period (90 days in this case). This feature is designed to protect Key Vault from accidental deletion of keys, secrets, and certificates.

The feature you should use to enforce retention period and accidental deletion is Purge Protection.



Purge Protection, when enabled, ensures that key vault objects cannot be permanently deleted until the retention period has passed. This feature provides an additional layer of protection to prevent the permanent deletion of key vault objects during the retention period.

upvoted 1 times

  **leviatas** 1 year, 9 months ago

Got this question in exam. 2023.09.25

upvoted 1 times

  **applepie** 1 year, 11 months ago



got this question today, go with the provided answer - 7/30/2023, score 895/1000

upvoted 3 times

  **NightshadeRC** 1 year, 11 months ago

Had this question in today's exam: 2023-07-26

upvoted 2 times

  **murat12345** 1 year, 11 months ago

If Box 1 is about having a retention period (meaning: you cant hard delete before the 90 days pass), what additional advantage does box 2 / purge protection give?

upvoted 1 times

  **basquiat** 1 year, 10 months ago



The first is related to key vault secrets and the second to the keyvault itself, from what I understood.

upvoted 1 times

  **Juhi\_Az204** 2 years, 3 months ago

Got this on 26March,2023

upvoted 1 times

  **Esward** 2 years, 5 months ago

Given answers are correct

<https://learn.microsoft.com/en-us/azure/key-vault/general/soft-delete-overview>

upvoted 1 times

  **Lucifer14** 2 years, 6 months ago

In 28-12-2022 exam

upvoted 1 times

  **xRiot007** 3 years ago

You use Soft delete to \*enable\* and Purge protection to \*enforce\*



upvoted 3 times

  **Eltooth** 3 years ago

Soft delete

Purge protection

upvoted 1 times

  **Rini100** 3 years, 2 months ago

Got this on 20 Apr 2022

upvoted 3 times

🗨️ 👤 **meoukg** 3 years, 3 months ago  
ot it on 03/2022, I chose Soft delete => Purge protection  
upvoted 4 times

🗨️ 👤 **francis6170** 4 years ago  
got this in the exam :)  
upvoted 7 times

🗨️ 👤 **UnknowMan** 4 years, 1 month ago  
The answer is correct.  
upvoted 2 times

🗨️ 👤 **glam** 4 years, 1 month ago  
correct  
upvoted 3 times

You provide an Azure API Management managed web service to clients. The back-end web service implements HTTP Strict Transport Security (HSTS).

Every request to the backend service must include a valid HTTP authorization header.

You need to configure the Azure API Management instance with an authentication policy.


Which two policies can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Basic Authentication
- B. Digest Authentication
- C. Certificate Authentication
- D. OAuth Client Credential Grant

**Suggested Answer:** CD

Community vote distribution



**Nielson** Highly Voted 4 years, 4 months ago

A, C

<https://www.examttopics.com/exams/microsoft/az-203/view/15/>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>

upvoted 104 times

**BrettusMaximus** 4 years, 2 months ago

Must be A,C As the API documentation only allows 3 options. It states: >>>>

Authentication policies

Authenticate with Basic - Authenticate with a backend service using Basic authentication.

Authenticate with client certificate - Authenticate with a backend service using client certificates.

Authenticate with managed identity - Authenticate with the managed identity for the API Management service.

upvoted 22 times

**BrettusMaximus** 4 years, 2 months ago

It can't be D as we need to authenticate to the backend and not from client.

upvoted 9 times

**1CY1** 1 year ago

While it does work, I thought Basic Auth was deprecated in a lot of systems now because it's not considered sufficiently secure. -> C and D

upvoted 1 times

**1CY1** 11 months ago

Because MS does not recommend Basic I might end up going with C and D.

upvoted 1 times

**fadikh** 4 years, 3 months ago

but client certificate does not use the authorization header

upvoted 6 times

**vb3d** 4 years, 3 months ago

I think this means correct answer is A and D

upvoted 8 times

**vb3d** 4 years, 3 months ago

Client certificate does use the x-arr-clientcert header

As Kitkit pointed out below, check this link

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>.



upvoted 11 times

**hobob** Highly Voted 4 years, 3 months ago

A, D

The question states that the solution MUST use the HTTP authorization header. The only two options that do are Basic Authentication (where the header would be "Authorization: Basic <auth-base64-encoded-string> or "Bearer <bearer-token-string>).



upvoted 40 times

  **pmsiva** 4 years, 3 months ago

AD is correct. Certificate authentication will not send http auth header

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies#Basic>

upvoted 6 times



  **syfool** 3 years, 10 months ago

D is not the answer, as that is not even an AUTHENTICATION POLICY!!!

The link below lists 3 authentication policies:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>

upvoted 6 times

  **MiraA** 3 years, 9 months ago

It seems "OAuth Client Credential Grant" exists and could be considered as the specific type of "Authenticate with managed identity" which sends the authorization header?

<https://oauth.net/2/grant-types/client-credentials/>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>

upvoted 5 times

  **2f17d52** Most Recent 2 months ago

**Selected Answer: AD**

From Chatgpt:

A. Basic Authentication

APIM can add a Basic Authorization header by encoding a username and password.

D. OAuth Client Credential Grant

APIM can be configured to obtain a token using client credentials flow and include a Bearer token in the Authorization header.

B. Digest Authentication

Not supported by Azure API Management.

C. Certificate Authentication

While APIM supports client certificate authentication, this is not implemented as an Authorization header, but at the TLS layer.

Does not meet the requirement of adding an HTTP Authorization header.



upvoted 1 times

  **rodso** 4 months, 3 weeks ago

**Selected Answer: CD**

Basic has security issues

upvoted 1 times

  **Vichu\_1607** 8 months, 3 weeks ago

**Selected Answer: AD**

A. Basic Authentication D. OAuth Client Credential Grant

Basic Authentication and OAuth Client Credential Grant are two policies that can be used to include a valid HTTP authorization header in every request to the backend service.

Basic Authentication sends a Base64-encoded string that contains a username and password for authentication.

OAuth Client Credential Grant is used when the client is requesting access to the protected resources under its control (i.e., it's own credentials), or those of another resource owner which have been previously arranged with the authorization server (the client is not impersonating the resource owner).

Digest Authentication and Certificate Authentication are not typically used for sending an HTTP authorization header.

upvoted 2 times

  **8ac3742** 10 months, 3 weeks ago

The certificate doesn't use an authorization header, it is mutual TLS(Transport Layer Security)

upvoted 1 times

🗳️ 👤 **Christian\_garcia\_martin** 11 months, 1 week ago

A and D are correct , about C Certificate Authentication is a valid form of authentication, it is not typically supported by Azure API Management for backend services

upvoted 1 times

🗳️ 👤 **FeriAZ** 1 year, 4 months ago

**Selected Answer: AD**

A. Basic Authentication: You can use this policy to add the Authorization header with Basic credentials.

D. OAuth Client Credential Grant: You can use this policy to acquire an OAuth token and add it to the Authorization header.

While Certificate Authentication (C) does indeed provide a method of authentication, it does not directly involve setting an HTTP Authorization header in the way Basic Authentication and OAuth Client Credential Grant do.

upvoted 3 times

🗳️ 👤 **raymond\_abcd** 1 year, 5 months ago

It is A and C

A, Basic Authentication: the Base064 encoded username and password are in the authentication header

C, Certificate base authentication: the client certificate is in the authentication header

upvoted 1 times

🗳️ 👤 **Ser1** 1 year, 6 months ago

**Selected Answer: AD**

.....

upvoted 1 times

🗳️ 👤 **ENGs** 1 year, 9 months ago

**Selected Answer: AD**

OAuth Client Credentials Grant (D), when used, will use an Authorization header being sent containing a bearer token.

Basic Authentication (A) will send an Authorization header.

--

Client Certificate (C) could be passed through an HTTP header but not necessarily the Authorization header.

upvoted 2 times

🗳️ 👤 **ck\_malviya** 1 year, 9 months ago

**Selected Answer: AD**

AD. The solution must use Authorization header

Basic auth and OAUTH user Bearer token auth header.

upvoted 1 times

🗳️ 👤 **Vladimir\_Gajinov** 1 year, 10 months ago

**Selected Answer: AD**

I would go with A and D.

A: allows you to require clients to include a valid username and password in the HTTP authorization header. While it's a simple method, it might not be the most secure option for modern applications, as the credentials are sent with each request in base64-encoded format.

D: OAuth Client Credential Grant is a flow where a client (in this case, your API Management instance) uses its own credentials (client ID and client secret) to obtain an access token from an OAuth authorization server. This access token can be included in the authorization header of requests to the backend service. This method is more secure than Basic Authentication and provides better control over authentication and authorization.

B: is a challenge-response mechanism that involves a server-provided nonce. It doesn't directly match the requirement for sending an HTTP authorization header with each request.

C: Certificate Authentication involves using client certificates to authenticate the client to the server. While this can be secure, it's not the same as sending an HTTP authorization header.

upvoted 1 times

🗳️ 👤 **applepie** 1 year, 11 months ago

got this question today, answer C,D without confidence - 7/30/2023, score 895/1000

upvoted 1 times

🗳️ 👤 **uncledana** 1 year, 11 months ago

I believe it is:

A - Sets the HTTP Authorization header to a corresponding value in the policy request

D - Authenticates to the backend with a valid HTTP authorization header

<https://learn.microsoft.com/en-us/azure/api-management/authentication-managed-identity-policy>

upvoted 1 times

  **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 1 times

  **minYoongi** 2 years, 2 months ago

AD is correct. <https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad#configure-a-jwt-validation-policy-to-pre-authorize-requests>

The following example policy, when added to the <inbound> policy section, checks the value of the audience claim in an access token obtained from Azure AD that is presented in the Authorization header. Straight from the documentation

upvoted 2 times



## DRAG DROP -

You are developing an ASP.NET Core website that can be used to manage photographs which are stored in Azure Blob Storage containers.

Users of the website authenticate by using their Azure Active Directory (Azure AD) credentials.

You implement role-based access control (RBAC) role permissions on the containers that store photographs. You assign users to RBAC roles.

You need to configure the website's Azure AD Application so that user's permissions can be used with the Azure Blob containers.

How should you configure the application? To answer, drag the appropriate setting to the correct location. Each setting can be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

**Settings**

client\_id

profile

delegated

application

user\_impersonation

**Answer Area**

API	Permission	Type
Azure Storage	Setting	Setting
Microsoft Graph	User.Read	Setting

**Suggested Answer:****Settings**

client\_id

profile

delegated

application

user\_impersonation

**Answer Area**

API	Permission	Type
Azure Storage	user_impersonation	delegated
Microsoft Graph	User.Read	delegated

Box 1: user\_impersonation -

Box 2: delegated -

Example:

1. Select the API permissions section
2. Click the Add a permission button and then:  
Ensure that the My APIs tab is selected
3. In the list of APIs, select the API TodoListService-aspnetcore.
4. In the Delegated permissions section, ensure that the right permissions are checked: user\_impersonation.
5. Select the Add permissions button.

Box 3: delegated -

Example -

1. Select the API permissions section
2. Click the Add a permission button and then,  
Ensure that the Microsoft APIs tab is selected
3. In the Commonly used Microsoft APIs section, click on Microsoft Graph
4. In the Delegated permissions section, ensure that the right permissions are checked: User.Read. Use the search box if necessary.
5. Select the Add permissions button

Reference:

<https://docs.microsoft.com/en-us/samples/azure-samples/active-directory-dotnet-webapp-webapi-openidconnect-aspnetcore/calling-a-web-api-in-an-aspnet-core-web-application-using-azure-ad/>

🗄️ 👤 **TonyMel** 1 year, 3 months ago  
correct, in 2023Mar24, score: 904/1000  
upvoted 12 times

🗄️ 👤 **mlantonis** Highly Voted 3 years, 1 month ago

Box 1: user\_impersonation

The built-in user\_impersonation scope indicates that the token is being requested on behalf of the user. Azure Storage exposes a single delegation scope named user\_impersonation that permits applications to take any action allowed by the user.

Box 2: delegated

Box 3: delegated

Reference:

<https://stackoverflow.com/questions/31404128/azure-ad-app-application-permissions-vs-delegated-permissions>

<https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-permissions-and-consent>

<https://docs.microsoft.com/en-us/azure/storage/common/storage-auth-aad-app?tabs=dotnet>

<https://docs.microsoft.com/en-us/rest/api/storageservices/authorize-with-azure-active-directory>

upvoted 38 times

🗄️ 👤 **p2006** Most Recent 9 months, 2 weeks ago

<https://learn.microsoft.com/en-us/information-protection/develop/concept-api-permissions#delegated-permissions>

upvoted 2 times

🗄️ 👤 **Dianahu** 11 months, 1 week ago

<https://www.andrew-best.com/posts/please-sir-can-i-have-some-auth/> first image

upvoted 1 times

🗄️ 👤 **Dats1987** 1 year, 2 months ago

That's Correct.

upvoted 1 times

🗄️ 👤 **Priya0703** 1 year, 4 months ago

On exam today 20-02-2023

upvoted 1 times

🗄️ 👤 **Esward** 1 year, 5 months ago

Given answers are correct

<https://stackoverflow.com/questions/31404128/azure-ad-app-application-permissions-vs-delegated-permissions>

upvoted 1 times

🗄️ 👤 **serpevi** 1 year, 9 months ago

Got this in 09/22 , went with the most voted answers, score 927.

upvoted 6 times

🗄️ 👤 **Eltooth** 2 years ago

User

Delegated

Delegated

upvoted 3 times

🗄️ 👤 **petitbilly** 2 years, 3 months ago

Got it in exam 03/22

upvoted 3 times

🗄️ 👤 **monaindia** 2 years ago

what answer you have chosen

upvoted 1 times

🗄️ 👤 **edengoforit** 2 years, 7 months ago

Application Permissions: Your application needs to access the web API directly as itself (no user context). This type of permission requires administrator consent and is also not available for native client applications.

Delegation Permissions: Your application needs to access the web API as the signed-in user, but with access limited by the selected permission. This type of permission can be granted by a user unless the permission is configured as requiring administrator consent.

upvoted 4 times

🗨️ 👤 **lugospod** 2 years, 5 months ago

Nope, it explicitly states that user is given RBAC permissions, and that each users permissions need to be used to access storage. So Application is not an option.

upvoted 1 times

🗨️ 👤 **jungaster** 3 years, 1 month ago

the answer is correct.

upvoted 1 times

🗨️ 👤 **roybunt** 3 years, 1 month ago

The answer is correct. For granting registered app permissions to Azure Storage, find in reference: <https://docs.microsoft.com/en-us/azure/storage/common/storage-auth-aad-app?tabs=dotnet#grant-your-registered-app-permissions-to-azure-storage>

upvoted 4 times

🗨️ 👤 **glam** 3 years, 1 month ago

correct

upvoted 2 times

🗨️ 👤 **kwaazaar** 3 years, 2 months ago

why Graph API?

upvoted 4 times

🗨️ 👤 **Sachini** 2 years, 4 months ago

"Permissions are granted to Microsoft Graph automatically when you first register your app with Azure AD."

upvoted 2 times

## HOTSPOT -

You are developing an ASP.NET Core app that includes feature flags which are managed by Azure App Configuration. You create an Azure App Configuration store named AppFeatureFlagStore that contains a feature flag named Export.

You need to update the app to meet the following requirements:

- ⇒ Use the Export feature in the app without requiring a restart of the app.
- ⇒ Validate users before users are allowed access to secure resources.
- ⇒ Permit users to access secure resources.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```

public void Configure(IApplicationBuilder app, IWebHostEnvironment env)
{
    if (env.IsDevelopment())
    {
        app.UseDeveloperExceptionPage();
    }
    else
    {
        app.UseExceptionHandler("/Error");
    }
}

app. [dropdown] ();
    UseAuthentication
    UseStaticFiles
    UseSession
    UseCookiePolicy

app. [dropdown] ();
    UseAuthorization
    UseHttpsRedirection
    UseSession
    UseCookiePolicy

app. [dropdown] ();
    UseAzureAppConfiguration
    UseRequestLocalization
    UseCors
    UseStaticFiles

app.UseEndpoint(endpoints =>
{
    endpoints.MapRazorPages();
});
}

```

## Answer Area

Suggested Answer:

```
public void Configure(IApplicationBuilder app, IWebHostEnvironment env)
{
    if (env.IsDevelopment())
    {
        app.UseDeveloperExceptionPage();
    }
    else
    {
        app.UseExceptionHandler("/Error");
    }

    app. [dropdown] ();
    UseAuthentication
    UseStaticFiles
    UseSession
    UseCookiePolicy

    app. [dropdown] ();
    UseAuthorization
    UseHttpsRedirection
    UseSession
    UseCookiePolicy

    app. [dropdown] ();
    UseAzureAppConfiguration
    UseRequestLocalization
    UseCors
    UseStaticFiles

    app.UseEndpoint(endpoints =>
    {
        endpoints.MapRazorPages();
    });
}
```

Box 1: UseAuthentication -

Need to validate users before users are allowed access to secure resources.

UseAuthentication adds the AuthenticationMiddleware to the specified IApplicationBuilder, which enables authentication capabilities.

Box 2: UseAuthorization -

Need to permit users to access secure resources.

UseAuthorization adds the AuthorizationMiddleware to the specified IApplicationBuilder, which enables authorization capabilities.

Box 3: UseStaticFiles -

Need to use the Export feature in the app without requiring a restart of the app.

UseStaticFiles enables static file serving for the current request path

Reference:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.aspnetcore.builder.iapplicationbuilder?view=aspnetcore-5.0>

 **hobob** Highly Voted 4 years, 3 months ago

UseAuthentication, Use Authorization, UseAzureAppConfiguration.

The last option allows dynamic configuration to be reloaded without an app restart.

<https://docs.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core5x>

upvoted 166 times

 **Pooochie** 4 years, 3 months ago

According to the site provided by you it should be

1. UseAzureAppConfiguration

2. UseAuthentication



3. Use Authorization

upvoted 13 times

 **xRiot007** 3 years ago

It is irrelevant in this case, the boxes only have ONE correct option each and those are, box1: authentication, box2: authorization, box3: app config.

upvoted 2 times

  **hobob** 4 years, 3 months ago



And that's what I said. Just in a different order.

upvoted 9 times

  **avanthasiriwardana** 4 years, 3 months ago

Order matters



upvoted 12 times

  **Kerrits** 4 years, 2 months ago

The UseAzureAppConfiguration is not available in the first list of options, only the last list of options.

Also, for configuring the app the order only matters sometimes. It doesn't in this case.

upvoted 7 times

  **OPT\_001122** 2 years, 10 months ago

i have doubt regarding UseAuthentication

UseAuthentication - this is not in the code snippet in the above link

```
public void Configure(IApplicationBuilder app, IWebHostEnvironment env)
{
    .....
    app.UseAzureAppConfiguration();

    app.UseHttpsRedirection();

    app.UseStaticFiles();

    app.UseRouting();



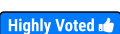
    app.UseAuthorization();
}
```

upvoted 1 times

  **rdemontis** 4 years, 3 months ago

correct, you're right!

upvoted 4 times

  **mlantonis**  4 years, 1 month ago

Box 1: UseAuthentication

Need to validate users before users are allowed access to secure resources.

UseAuthentication adds the AuthenticationMiddleware to the specified IApplicationBuilder, which enables authentication capabilities. To verify users, we need to use the Authentication middleware.

Box 2: UseAuthorization

Need to permit users to access secure resources.

UseAuthorization adds the AuthorizationMiddleware to the specified IApplicationBuilder, which enables authorization capabilities. To authorize users to access resources , we need to use the Authorization middleware.

Box 3: UseAzureAppConfiguration

Adding the UseAzureAppConfiguration middleware to allow the configuration settings registered for refresh to be updated while the ASP.NET Core web app continues to receive requests. For using feature flags, you need to make use of the Azure App Configuration service. To ensure the configuration settings are refreshed without the need to restart the web app, you can use the middleware of UseAzureAppConfiguration();

upvoted 60 times

  **mlantonis** 4 years, 1 month ago

Reference:

<https://docs.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core5x>

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.aspnetcore.builder.iapplicationbuilder?view=aspnetcore-5.0>

upvoted 12 times

  **ciamp**  10 months, 2 weeks ago

the middleware get declare in specific sequence so

UseAuthentication

UseAuthorization

UseAzureAppConfiguration


upvoted 1 times

  **raymond\_abcd** 1 year, 5 months ago

Answer is: UseAuthentication, UseAuthorization, UseAzureAppConfiguration. You need to use the feature flags for the requirement to enable the export feature. For that you have to call UseAzureAppConfiguration. See: <https://learn.microsoft.com/en-us/azure/azure-app-configuration/use-feature-flags-dotnet-core?tabs=core6x>



The UseStaticFiles should normally be called first before authentication, but it does not meet any requirement to choose this as an option.

upvoted 1 times

  **HimGarg** 1 year, 8 months ago

Why no one from ExamTopics bother to correct the answers?

upvoted 4 times

  **Ciupaz** 1 year, 5 months ago

This is the real mystery.

upvoted 1 times

  **gmishra88** 2 years, 8 months ago

And Microsoft claims this exam does not require specific knowledge of .Net SDK.

upvoted 6 times

  **macobuzi** 1 year, 10 months ago

I doubt it ever stated that

upvoted 1 times

  **tushar1812** 2 years, 11 months ago

Drop down 1 UseAuthentication

Drop down 2 Use Authorization

Drop down 3 UseAzureAppConfiguration.

The last option allows dynamic configuration to be reloaded without an app restart.

Refer

<https://docs.microsoft.com/en-us/azure/azure-app-configuration/enable-dynamic-configuration-aspnet-core?tabs=core5x>

upvoted 2 times

  **Eltooth** 3 years ago

Authentication

Authorisation

Appconfig

upvoted 3 times

  **tramlong888** 3 years, 6 months ago

the correct answer should be UseAuthentication, Use Authorization, UseAzureAppConfiguration.

upvoted 3 times

  **doudou123** 3 years, 9 months ago

according Udemy,

1. UseAuthentication

- 2. Use Authorization
  - 3. Use AzureAppConfiguration
- upvoted 5 times

  **burakcive** 3 years ago



Which course on udemy?  
upvoted 1 times

  **kondapaturi** 4 years ago

app.UseAuthentication(); - To verify users ,we need to use the Authentication middleware.  
app.UseAuthorization(); -To authorize users to access resources , we need to use the Authorization middleware.  
app.UseAzureAppConfiguration(); -For using feature flags, you need to make use of the Azure App Configuration service. To ensure the configuration settings are refreshed without the need to restart the web app, you can use the middleware of UseAzureAppConfiguration();  
upvoted 2 times

  **UnknowMan** 4 years, 1 month ago

- 1. UseAuthentication, -> Login
  - 2. Use Authorization, -> Authorization
  - 3. UseAzureAppConfiguration. -> Dynamic App settings
- upvoted 3 times

  **glam** 4 years, 1 month ago

UseAuthentication, Use Authorization, UseAzureAppConfiguration.  
upvoted 3 times

  **pavan555manjunath** 4 years, 1 month ago

Any one please confirm the answer is correct

- 1. UseAuthentication,
- 2. Use Authorization,
- 3. UseAzureAppConfiguration.

upvoted 6 times



You have an application that includes an Azure Web app and several Azure Function apps. Application secrets including connection strings and certificates are stored in Azure Key Vault.

Secrets must not be stored in the application or application runtime environment. Changes to Azure Active Directory (Azure AD) must be minimized.

You need to design the approach to loading application secrets.

What should you do?

- A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.
- B. Create a single Azure AD Service Principal with permission to access Key Vault and use a client secret from within the App Services to access Key Vault.
- C. Create a system assigned Managed Identity in each App Service with permission to access Key Vault.
- D. Create an Azure AD Service Principal with Permissions to access Key Vault for each App Service and use a certificate from within the App Services to access Key Vault.

**Suggested Answer: C**


Use Key Vault references for App Service and Azure Functions.

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.

Reference:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>

*Community vote distribution*



**YellowSky002** 4 years, 3 months ago

I think the answer is A. Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity.

Because we have more than one App (Web App and other Function Apps) , So we agree it is going to be a managed identity but should I create one for each app or one for all apps?

If I create system MI then there should be one for each App.

If I create user MI then I can re-use it for any App I want with minimum change to AD

upvoted 118 times

**Basu525** 4 years, 3 months ago

Yes exactly. The question emphasizes on having multiple resources having to access Key vault. So using User Assigned MI is more pragmatic.

upvoted 4 times

**SlavMar** 4 years ago

Well from security standpoint you could assign different secrets to different parts of application (if they are using system managed identity)

soi if one component is compromised - not all secrets are available to hackers

upvoted 5 times

**midhy** 4 years, 3 months ago

User assigned cannot be used Look at the link provided.

upvoted 4 times

**vb3d** 4 years, 3 months ago

I think that note in the link is only for the usage of key vault references, which is a new thing in key vaults. I think a user assigned identity should still have access to the key vault.

upvoted 4 times

**Zsolt72** 4 years, 3 months ago



Yes, key vault reference is an App Configuration feature referencing a secret from key vault.

upvoted 1 times



**Tom87** 4 years, 2 months ago

That's true.



<https://kasunkodagoda.com/2019/06/09/using-user-assigned-managed-identity-to-access-azure-key-vault-from-azure-app-service/>  
upvoted 1 times

  **vb3d** 4 years, 3 months ago

I just tried this, I was able to give access to a user assigned managed identity to a key vault.  
upvoted 8 times

  **vtomy** 3 years, 12 months ago



User assigned MI is not supported for keyvault.  
System assigned MI is correct.  
upvoted 6 times

  **Philli** 3 years, 6 months ago

I don't read this on that page t.b.h.  
upvoted 1 times




  **zzt** 3 years, 11 months ago

This is not true: see: <https://kasunkodagoda.com/2019/06/09/using-user-assigned-managed-identity-to-access-azure-key-vault-from-azure-app-service/>  
upvoted 6 times

  **ferut** 4 years, 1 month ago

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>  
The lifecycle of user-assigned managed identity is manually managed by the user and can be used (not assigned) to several services.  
The lifecycle of the system-assigned managed identity is tightly coupled with the service, it's assigned to a service. When the service is deleted, the identity will also be gone.

To minimize changes to AD, user-assigned is more suitable.  
upvoted 8 times

  **mlantonis** Highly Voted  4 years, 1 month ago

Correct Answer: C

Key Vault references currently only support system-assigned managed identities. User-assigned identities cannot be used.



Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types>  
<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>  
upvoted 41 times


  **altafpatel1984** 3 years, 7 months ago

Don't misguide people. In the same link you give, this statement exists "Some apps need to reference secrets at creation time, when a system-assigned identity would not yet be available. In these cases, a user-assigned identity can be created and given access to the vault in advance."

Your given reference: <https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>  
upvoted 10 times

  **zzt** 3 years, 11 months ago

Not correct: see link <https://kasunkodagoda.com/2019/06/09/using-user-assigned-managed-identity-to-access-azure-key-vault-from-azure-app-service/>  
upvoted 5 times

  **eX900** 2 years, 8 months ago

Granting your app access to Key Vault  
2.Create a managed identity for your application.  
"Key Vault references will use the app's system assigned identity by default, but you can specify a user-assigned identity."

Your second reference:  
<https://docs.microsoft.com/en-us/azure/app-service/app-service-key-vault-references>.

So, the answer is : A  
upvoted 4 times

🗳️ 👤 **TakumaK** 3 years, 12 months ago

You are wrong. your second link OBVIOUSLY says

"Key Vault references will use the app's system assigned identity by default, but you can specify a user-assigned identity."

upvoted 19 times

🗳️ 👤 **Vichu\_1607** Most Recent 8 months, 3 weeks ago

Selected Answer: A

Using a user-assigned Managed Identity is a good approach in this scenario. Managed identities provide an identity for applications to use when connecting to resources that support Azure Active Directory (Azure AD) authentication.

A user-assigned managed identity is created as a standalone Azure resource. After it's created, the identity can be assigned to one or more Azure service instances. The identity is managed by Azure AD and does not require a password.

By using a single user-assigned Managed Identity across all your App Services, you can manage access to the Key Vault centrally, minimizing changes to Azure AD.

upvoted 1 times

🗳️ 👤 **prabhjot** 10 months, 3 weeks ago

Ans is C - Secrets must not be stored in the application or application runtime environment.

upvoted 1 times

🗳️ 👤 **jobolesoniha1** 1 year, 3 months ago

Selected Answer: A

Key vault references use the app's system-assigned identity by default, but you can specify a user-assigned identity.

Even for KeyVault References - "Key vault references use the app's system-assigned identity by default, but you can specify a user-assigned identity."

upvoted 1 times

🗳️ 👤 **Zhenyuan\_Zhang** 1 year, 4 months ago

Selected Answer: A

I think the answer is A.

<https://learn.microsoft.com/en-us/entra/identity/managed-identities-azure-resources/overview#management-identity-types>  
<https://docs.microsoft.com/en-us%20azure/app-service/app-service-key-vault-references>

upvoted 1 times

🗳️ 👤 **junkz** 1 year, 8 months ago

there is a saying in my country: one fool throws a rock in the pond, and 10 good people do their best to pull it out. this is exactly how i feel looking at all of us trying to "guess" what "minimizing AD changes" really means in this context...

upvoted 4 times

🗳️ 👤 **drkpns1** 1 year, 10 months ago

I think in the past, you cannot use the user-assigned identity to access to the Azure Key-Vault. But it is possible now: "Key vault references use the app's system-assigned identity by default, but you can specify a user-assigned identity."

<https://learn.microsoft.com/en-us/azure/app-service/app-service-key-vault-references?tabs=azure-cli>

So I think the answer should be A.

upvoted 2 times

🗳️ 👤 **Dianahu** 1 year, 11 months ago

if I want to add an user-assigned managed identity I need to create it so I wouldnt say that it is minimum change to AD and it is not a good pattern to have the same identity for all of them

upvoted 1 times

🗳️ 👤 **AzSimuz** 2 years, 2 months ago

Managed identities are of the following two types:

System-assigned -- this is activated directly on the Azure cloud service. For example, you can create a managed identity for your web app that will be automatically deleted once the web app instance is deleted

User-assigned -- this is a stand-alone Azure resource that is created independent of the actual service instance. You can assign this identifier to one or more Azure resources simultaneously. As an example, you can set this identity to both an Azure Web App and an Azure Function at the same time

- <https://www.loginradius.com/blog/engineering/guest-post/using-azure-key-vault-with-an-azure-web-app-in-c-sharp/>

upvoted 2 times

🗳️ 👤 **maqsoodshah** 2 years, 3 months ago

Correct Answer: A

System-assigned -- this is activated directly on the Azure cloud service. For example, you can create a managed identity for your web app that will be automatically deleted once the web app instance is deleted

User-assigned -- this is a stand-alone Azure resource that is created independent of the actual service instance. You can assign this identifier to one or more Azure resources simultaneously. As an example, you can set this identity to both an Azure Web App and an Azure Function at the same time  
upvoted 1 times

🗳️ 👤 **JamieS** 2 years, 3 months ago

I would have thought it was system assigned identity, but it says we have multiple Azure Function apps, and one of the requirements is to minimise changes to AD.

Only the user identity has a one to many mapping (<https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview#managed-identity-types>).

upvoted 3 times

🗳️ 👤 **manopeydakon** 2 years, 4 months ago

Key Vault references will use the app's system assigned identity by default, but you can specify a user-assigned identity.

User-assigned managed identity Can be shared!

The same user-assigned managed identity can be associated with more than one Azure resource.

So A and C both works but less work with A in this case.

upvoted 1 times

🗳️ 👤 **Esward** 2 years, 5 months ago

User assigned Identity

<https://thecodeblogger.com/2020/06/13/user-assigned-managed-identity-with-azure-key-vault/>

upvoted 1 times

🗳️ 👤 **Esward** 2 years, 5 months ago

Also refer

<https://yourazurecoach.com/2022/02/04/use-app-service-key-vault-references-with-user-assigned-identities/>

upvoted 1 times

🗳️ 👤 **Esward** 2 years, 5 months ago

<https://stackoverflow.com/questions/55101525/how-can-i-give-access-to-key-vault-to-a-user-assigned-identity>

upvoted 1 times

🗳️ 👤 **daothanhien** 2 years, 5 months ago

**Selected Answer: C**

C.

Because user assigned will take more codes and configurations for the key.

System assigned for each app service more security and noneed to have any configuration and the same code

upvoted 4 times

🗳️ 👤 **imanonion** 2 years, 5 months ago

**Selected Answer: C**

By creating a system-assigned Managed Identity for each App Service, you can avoid storing secrets in the application or runtime environment, as the Managed Identity provides a way for the App Service to authenticate to Key Vault without the need for hard-coded credentials. Additionally, using a separate Managed Identity for each App Service allows for fine-grained access control to the secrets stored in Key Vault, as you can assign permissions to specific identities rather than granting broad permissions to a single identity.

upvoted 1 times


🗳️ 👤 **imanonion** 2 years, 5 months ago

Option A, "Create a single user-assigned Managed Identity with permission to access Key Vault and configure each App Service to use that Managed Identity" also can be a good solution, But it can't provide fine-grained access control as all App Service using the same identity.

Option B "Create a single Azure AD Service Principal with permission to access Key Vault and use a client secret from within the App Services to access Key Vault" would not be the recommended approach, as it would involve storing the client secret in the App Services, which is a violation of the requirement to avoid storing secrets in the application or runtime environment.

Option D "Create an Azure AD Service Principal with Permissions to access Key Vault for each App Service and use a certificate from within the App Services to access Key Vault" is also a valid approach but it is more difficult to manage the certificate in each App Service and need more code to handle the certificate in the code.

upvoted 1 times

  **warchoon** 2 years, 4 months ago

User-assigned Managed Identity is an additional resource. System-assigned is not

upvoted 2 times

  **oceane0316** 2 years, 6 months ago

The correct answer is A.

Source : <https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/managed-identity-best-practice-recommendations#using-user-assigned-identities-to-reduce-administration>

Read the article but in summary "When a user-assigned identity is associated with the four virtual machines, only two role assignments are required, compared to eight with system-assigned identities. If the virtual machines' identity requires more role assignments, they'll be granted to all the resources associated with this identity."

The key here is "Changes to Azure Active Directory (Azure AD) must be minimized" and the paragraph is called "Using user-assigned identities to reduce administration".

So A is the correct answer because of the "change must be minimized". C would work but it means you have much more changes.

upvoted 3 times

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 are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised. You need to store the intake forms according to the requirements.

Solution:

1. Create an Azure Key Vault key named skey.
2. Encrypt the intake forms using the public key portion of skey.
3. Store the encrypted data in Azure Blob storage.


Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: A**

Community vote distribution



**pac1311** Highly Voted 4 years, 4 months ago

Correct!

upvoted 37 times

**bingohah** Highly Voted 4 years, 3 months ago

The answer is right, but not good. Because we should use symmetric key to encrypt file, instead of public key. The best practise is: encrypt file by symmetric key, and the symmetric key is protected by public key.

upvoted 29 times

**vb3d** 4 years, 3 months ago

Not sure what do you mean, but check the link below.

Symmetric key means the same key is used for encryption and decryption. When you use an Asymmetric key, you encrypt with the public key and decrypt with a secret key, which is what the solution is talking about and is better.

<https://www.ssl2buy.com/wiki/symmetric-vs-asymmetric-encryption-what-are-differences#:~:text=Symmetric%20encryption%20uses%20a%20single,and%20decrypt%20messages%20when%20communicating.>

upvoted 11 times

**aarrtturas** 4 years, 1 month ago

We should use symmetric key not asymmetric key. Asymmetric key is used for small amount's of data to encrypt. Usually asymmetric key encrypts symmetric key, and symmetric key encrypts data

upvoted 5 times

**fearoffree** 3 years, 4 months ago

The asymmetric key is used to sign the intake so it could not be compromised!

upvoted 1 times

**J\_obando13** Most Recent 1 month, 2 weeks ago

**Selected Answer: A**

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 2 times

**2f17d52** 2 months ago

**Selected Answer: B**

From Chatgpt:

Problem: Directly Encrypting the Forms with the Key Vault RSA Key

Azure Key Vault RSA keys are asymmetric and typically support encrypting only small payloads (e.g., up to 4 KB).

Scanned medical forms are typically large files (hundreds of KB to several MB).

Trying to encrypt these large files directly with the RSA public key stored in Key Vault will:

Fail due to size limits, or

Be extremely inefficient and non-scalable.

upvoted 1 times

🗳️ 👤 **Vichu\_1607** 8 months, 3 weeks ago

**Selected Answer: A**

The solution does meet the goal. Azure Key Vault is a cloud service for securely storing and accessing secrets. A secret is anything that you want to tightly control access to, such as API keys, passwords, certificates, or cryptographic keys.

In this case, the intake forms are encrypted using the public key of a key stored in Azure Key Vault (skey). This means that even if the data is downloaded by a third party, they would not be able to decrypt and read the data without access to the corresponding private key.

Storing the encrypted data in Azure Blob Storage is a common and secure method for storing large amounts of unstructured data in the cloud.

upvoted 1 times

🗳️ 👤 **obllew** 9 months ago

It's poorly worded - is the 3rd party an authorized party? By "compromised" do they mean confidential info would be revealed? Are they saying docs should be encrypted at rest such that if someone gains access to the account, patient data would not be leaked? In which case A meets that goal, because only the website may securely access the doc by accessing the private key from Key Vault to decrypt them.

upvoted 2 times

🗳️ 👤 **CafeLago** 10 months, 2 weeks ago

Shouldn't this be WORM with immutable policies and NOT CMK or client-side encryption? There is no requirement for Confidentiality, only Integrity. Without mention of sharing the key with the third-party, it cannot read the intake forms and so the solution does not meet the goal.

upvoted 1 times

🗳️ 👤 **richardc7** 1 year, 9 months ago

**Selected Answer: A**

Got 29/09/2023

upvoted 3 times

🗳️ 👤 **notedo** 1 year, 9 months ago

**Selected Answer: A**

Correct

upvoted 1 times

🗳️ 👤 **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 4 times

🗳️ 👤 **Saluk\_DE** 2 years, 3 months ago

**Selected Answer: A**

Question was in Exam 2023-03-30

upvoted 2 times

🗳️ 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Chosen yes. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 1 times

🗳️ 👤 **Bear\_Polar** 2 years, 7 months ago

**Selected Answer: A**

AZ-900, is that you? :)))

upvoted 1 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: A**

A. Yes

upvoted 2 times

🗳️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: A**

A. Yes

upvoted 2 times

☐  **coffecold** 2 years, 8 months ago

The answer is B (NO)

This is a business case of digital signing.

In digital signing the sender's private key encrypts the data -- this is the digital signature -- and the receiver uses the public key to decrypt it and verify it matches the attachment.

"2. Encrypt the intake forms using the public key portion of skey."

must be

"2. Encrypt the intake forms using the private key portion of skey."

Digital signing is opposed to encryption : Encryption uses a public key to ensure the ciphertext cannot be deciphered by anyone but the authorized recipient with the private key.

upvoted 4 times

☐  **gmishra88** 2 years, 8 months ago

Client side encryption is a valid method to do this. Asymmetric key can be used (RSA in keyvault). If using Asymmetric key then client side encryption should be done with the public key because it can only be decrypted with a private key (which other people do not have). One will encrypt with private key only as a signature to prove she has the private key so that it can be verified with the public key.

But still a difficult question in a certification exam with a time limit if you used their SDK to do the encryption and decryption.

upvoted 1 times

☐  **gmishra88** 2 years, 8 months ago

And, I will always use a symmetric key for client side encryption. But god knows what Microsoft want as an answer and what the question creator experienced in life.

upvoted 1 times



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 are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised. You need to store the intake forms according to the requirements.

Solution:

1. Create an Azure Cosmos DB database with Storage Service Encryption enabled.
2. Store the intake forms in the Azure Cosmos DB database.


Does the solution meet the goal?

- A. Yes
- B. No

**Suggested Answer: B**

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

Community vote distribution



**cbn** Highly Voted 3 years, 10 months ago

I think the rationale is Azure Cosmos is not the best place to store blob data, like scanned patient intake forms.

<https://stackoverflow.com/questions/49343675/how-can-i-store-images-in-azure-cosmos-db>  
upvoted 31 times

**iamstudying** 2 years, 9 months ago

Although it's not ideal, it is still a valid solution (sort of)... media files are stored in blob storage <https://docs.microsoft.com/en-us/azure/cosmos-db/database-encryption-at-rest>

The key thing to point here is that the solution requires encrypted files when downloaded, but the solution here is only encryption at rest  
upvoted 12 times

**gmishra88** 2 years, 2 months ago

One of those typical Microsoft questions that punishes people who know something. If you do not understand the question then you will answer No, correctly (don't store images in cosmos), but if you know encryption to the next level you think it is encrypted so it is Yes. But one more level if you think then server side encryption (at rest) is not enough if it is already downloaded.  
upvoted 5 times

**vizay** 1 month, 4 weeks ago

✓ While Stored in Cosmos DB (Encryption at Rest)

Azure automatically encrypts the data on disk.

Even if someone gains unauthorized access to the storage system itself, they cannot read the raw data without the proper keys.

✓ This is encryption at rest – data is protected while stored.

✗ After Downloading the File

Now imagine:

A doctor downloads a patient's form from the web app to their personal laptop.

That laptop doesn't encrypt files by default.

Now, the file is sitting unprotected on the laptop – not encrypted anymore.

If that laptop is lost or hacked, the sensitive patient info can be read.

upvoted 1 times

  **vizay** 1 month, 4 weeks ago

▮ What Happens Step-by-Step:

You encrypt the file yourself (using something like AES-256).

Now the file is protected, and only someone with the right decryption key can read it.

You upload the encrypted file to Azure Blob Storage.

Azure automatically applies Storage Service Encryption (SSE) too – a second layer of encryption at rest.

Later, when the file is downloaded, it's still encrypted from your side – because you encrypted it before uploading.

▮ The person who downloads it can't read it without the decryption key.

✓ This solves the problem of files becoming readable after download.

upvoted 1 times

  **pieronegri** 3 years, 10 months ago

agreed

upvoted 2 times



  **mlantonis** Highly Voted 3 years, 7 months ago

Correct Answer: B - No

If storage is encrypted by default but you need to encrypt the content so when downloaded, its useless to unauthorised party.



Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

upvoted 30 times

  **TMakki** 2 years, 4 months ago

Absolutely right. Data is encrypted at rest but when downloaded it is not.

upvoted 3 times

  **J\_obando13** Most Recent 1 month, 2 weeks ago

Selected Answer: B

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

  **jobolesoniha1** 9 months ago

Selected Answer: B

Azure Storage Service Encryption, also known as server-side encryption (SSE), is a form of transparent data encryption<sup>1</sup>. This means that data is automatically encrypted when it is written to Azure Storage, and automatically decrypted when it is read<sup>1</sup>.

upvoted 1 times

  **MysticalSam** 1 year, 6 months ago

This question was in today's exam on 10-June-2023


upvoted 3 times

  **Saluk\_DE** 1 year, 9 months ago

Selected Answer: B

Question was in Exam 2023-03-30

upvoted 1 times

  **sarmaria** 1 year, 9 months ago

Got this on 16/03/23. Chosen No. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

🗨️ **hubekpeter** 2 years, 1 month ago

**Selected Answer: B**

For a god sake. This is an encryption at rest, where you are encrypting the underlying storage doesn't matter if it's with a azure or customer managed keys instead of encrypting each document itself on an upper layer.

upvoted 1 times

🗨️ **OPT\_001122** 2 years, 1 month ago

**Selected Answer: B**

No - correct ans

upvoted 1 times

🗨️ **eX900** 2 years, 2 months ago

If I have any doubts about the question, I would check these users (who are the top commenters) to be sure:

- Eltooth
- mlantonis
- glam

Thank you guys. 🙏

upvoted 4 times

🗨️ **Ciupaz** 1 year ago

Good tip, thanks.

upvoted 1 times

🗨️ **TMakki** 2 years, 4 months ago

**Selected Answer: B**

Data is encrypted at rest but when downloaded it is not. So when downloaded by third party it is not encrypted so instead use an Azure Key vault and public key encryption

upvoted 3 times

🗨️ **Knightie** 2 years, 4 months ago

**Selected Answer: A**

It is encrypted at rest, but available (not encrypted) for third party to download the copy, so the third party can read, come on. regarding encryption on transfer, it's the job of https.. so that's the way for third party to get the scanned copy, right?

upvoted 2 times

🗨️ **alcideus** 2 years, 4 months ago

Then why did you put A as your answer lol

upvoted 3 times

🗨️ **xRiot007** 2 years, 6 months ago

The answer is No. Encrypting the storage itself only. If someone requests it, they will get plain data.

upvoted 3 times

🗨️ **Eltooth** 2 years, 6 months ago

**Selected Answer: B**

B is correct answer.

upvoted 3 times

🗨️ **petitbilly** 2 years, 9 months ago

Got it in exam 03/22

upvoted 2 times

🗨️ **altafpatel1984** 3 years, 1 month ago

I guess Cosmos DB is for semi-structured data, not for unstructured data. Here scanned copies of forms are pdf/image files which are unstructured files, hence they cannot be stored in CosmosDB. Instead should be stored in Blob Storage which also supports security for data in-transit and data at rest. And because of this B - No is correct.

upvoted 2 times

🗨️ **glam** 3 years, 7 months ago

B. No..

upvoted 5 times

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 are developing a medical records document management website. The website is used to store scanned copies of patient intake forms. If the stored intake forms are downloaded from storage by a third party, the contents of the forms must not be compromised. You need to store the intake forms according to the requirements.

Solution: Store the intake forms as Azure Key Vault secrets.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

Community vote distribution


B (100%)

  **clarionprogrammer** Highly Voted 3 years, 2 months ago

B. No

You'd definitely not want to store "data" in Azure Key Vault.

upvoted 30 times

  **mlantonis** Highly Voted 3 years, 1 month ago

Correct Answer: B - No

You'd definitely not want to store "data" in Azure Key Vault.

Instead use an Azure Key vault and public key encryption. Store the encrypted from in Azure Storage Blob storage.

upvoted 18 times

  **Ciupaz** Most Recent 7 months, 3 weeks ago

"Azure Blob Storage", not "Azure Storage Blob Storage".

upvoted 1 times

  **richardc7** 9 months ago

Selected Answer: B

Got 28/09/2023

No

upvoted 1 times

  **MysticalSam** 1 year ago

This question was in today's exam on 10-June-2023

upvoted 2 times

  **Saluk\_DE** 1 year, 3 months ago

Selected Answer: B

Question was in Exam 2023-03-30

upvoted 1 times

  **sarmaria** 1 year, 3 months ago

Got this on 16/03/23. Chosen No. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

  **NombreFalso** 1 year, 4 months ago

No, that's crazy.

upvoted 3 times

🗨️ 👤 **OPT\_001122** 1 year, 7 months ago

**Selected Answer: B**

B. No -blob can not be stored in azure key vault  
upvoted 1 times

🗨️ 👤 **Eltooth** 2 years ago

**Selected Answer: B**

B is correct answer.  
upvoted 5 times

🗨️ 👤 **Rini100** 2 years, 2 months ago

Got this on 20 Apr 2022  
upvoted 3 times

🗨️ 👤 **Freidrich** 2 years, 4 months ago

**Selected Answer: B**

The correct answer is B: No.  
upvoted 1 times

🗨️ 👤 **Kvm1** 3 years ago

B. No  
Technically possible.. but KeyValut is intended to store securely Application secrets, certificates etc. and not to store application data like the Intake forms and other documents.  
upvoted 7 times

🗨️ 👤 **glam** 3 years, 1 month ago

B. No...  
upvoted 3 times

🗨️ 👤 **jokergester** 3 years, 2 months ago

Technically possible but the documentation insists on using the Key Vault as storage for Application data or config not user data.  
upvoted 3 times

## HOTSPOT -

You plan to deploy a new application to a Linux virtual machine (VM) that is hosted in Azure.

The entire VM must be secured at rest by using industry-standard encryption technology to address organizational security and compliance requirements.

You need to configure Azure Disk Encryption for the VM.

How should you complete the Azure CLI commands? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
```

az  ▼ create \

vm
keyvault
keyvault key
vm encryption

```
--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True
```

az  ▼ create \

vm
keyvault
keyvault key
vm encryption

```
--vault-name $keyvault_name \
--name Name1 \
--protection software
```

az  ▼ create \

vm
keyvault
keyvault key
vm encryption

```
--resource-group $resourcegroup \
--name Name2 \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \
--data-disk-sizes-gb 5
```

az  ▼ enable\

vm
keyvault
keyvault key
vm encryption

```
--resource-group $resourcegroup \
--name Name2 \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key Name1 \
--volume-type  ▼
```

all
data
os

## Answer Area

```
az provider register -n Microsoft.KeyVault
resourcegroup="myResourceGroup"
az group create --name $resourcegroup --location westus
keyvault_name=myvaultname$RANDOM
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 create \

--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 create \

--vault-name $keyvault_name \
--name Name1 \
--protection software
```

Suggested Answer:

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 create \

--resource-group $resourcegroup \
--name Name2 \
--image Canonical:UbuntuServer:16.04-LTS:latest \
--admin-username azureuser \
--generate-ssh-keys \
--data-disk-sizes-gb 5
```

```
az 

|               |   |
|---------------|---|
|               | ▼ |
| vm            |   |
| keyvault      |   |
| keyvault key  |   |
| vm encryption |   |

 enable\

--resource-group $resourcegroup \
--name Name2 \
--disk-encryption-keyvault $keyvault_name \
--key-encryption-key Name1 \
--volume-type 

|      |   |
|------|---|
|      | ▼ |
| all  |   |
| data |   |
| os   |   |


```

Box 1: keyvault -

Create an Azure Key Vault with az keyvault create and enable the Key Vault for use with disk encryption. Specify a unique Key Vault name for keyvault\_name as follows: keyvault\_name=myvaultname\$RANDOM az keyvault create \

```
--name $keyvault_name \
--resource-group $resourcegroup \
--location eastus \
--enabled-for-disk-encryption True
```

Box 2: keyvault key -

The Azure platform needs to be granted access to request the cryptographic keys when the VM boots to decrypt the virtual disks. Create a cryptographic key in your Key Vault with az keyvault key create. The following example creates a key named myKey: az keyvault key create \

```
--vault-name $keyvault_name \
--name myKey \
--protection software
```

Box 3: vm -

Create a VM with az vm create. Only certain marketplace images support disk encryption. The following example creates a VM named myVM using an Ubuntu

16.04 LTS image:

```
az vm create \  
--resource-group $resourcegroup \  
--name myVM \  
--image Canonical:UbuntuServer:16.04-LTS:latest \  
--admin-username azureuser \  
--generate-ssh-keys \  

```

Box 4: vm encryption -

Encrypt your VM with az vm encryption enable:

```
az vm encryption enable \  
--resource-group $resourcegroup \  
--name myVM \  
--disk-encryption-keyvault $keyvault_name \  
--key-encryption-key myKey \  
--volume-type all  

```

Note: seems to an error in the question. Should have enable instead of create.

Box 5: all -

Encrypt both data and operating system.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart>

  **Tom87**  2 years, 8 months ago

The answer is correct.

Other possibilities doesn't make sense. Although it seems strange to me to create the VM when the question says that we already have one.

upvoted 60 times

  **Tom87** 2 years, 8 months ago

Ok, now I see that the question doesn't say that the VM is already created. Sorry for that :)

upvoted 8 times

  **glam**  2 years, 7 months ago

correct

upvoted 8 times

  **Juhi\_Az204**  9 months ago

Got this on 26March,2023

upvoted 5 times

  **sarmaria** 9 months, 1 week ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 3 times

  **Esward** 11 months, 1 week ago

VM

Keyvault

Keyvault key

VM Enryption

ALL

Refer: <https://learn.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart>

upvoted 4 times

  **manopeydakon** 10 months, 2 weeks ago

Look at the whole picture son! Answer is correct

upvoted 3 times



🗒️ 👤 **rolling\_potato\_** 12 months ago

Was on my exam today (03-01-2023)  
upvoted 6 times

🗒️ 👤 **Lucifer14** 1 year ago

In 28-12-2022 exam  
upvoted 2 times

🗒️ 👤 **OPT\_001122** 1 year, 1 month ago

Key vault

Key vault key

VM

VM encryption

All

upvoted 3 times

🗒️ 👤 **at\_sayali** 1 year, 1 month ago

The answer is correct.

Got this in 10/2022

upvoted 2 times

🗒️ 👤 **gmishra88** 1 year, 2 months ago

The documentation here does not mention the step to create the keyvault key.

<https://learn.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart>.

But it is okay

upvoted 1 times

🗒️ 👤 **serpevi** 1 year, 3 months ago

Got this in 09/22 , went with the most voted answers, score 927.

upvoted 2 times

🗒️ 👤 **Eltooth** 1 year, 6 months ago

Answer is correct.

Keyvault

Keyvault key

VM

VM encryption

All

upvoted 4 times

🗒️ 👤 **Baskman** 1 year, 9 months ago

Got this in the exam 03/22

upvoted 5 times

🗒️ 👤 **oescm** 1 year, 10 months ago

Got this one 02/2022. Went with the given answer

upvoted 5 times

🗒️ 👤 **lugospod** 1 year, 11 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

🗒️ 👤 **tramlong888** 2 years ago

The answer is correct.

upvoted 1 times

🗒️ 👤 **debanjan10** 2 years, 2 months ago

Order: keyvault > keyvault key > vm > vm encryption

upvoted 4 times

Your company is developing an Azure API hosted in Azure.

You need to implement authentication for the Azure API to access other Azure resources. You have the following requirements:

- ⇒ All API calls must be authenticated.
- ⇒ Callers to the API must not send credentials to the API.

Which authentication mechanism should you use?

- A. Basic
- B. Anonymous
- C. Managed identity
- D. Client certificate

**Suggested Answer: C**

Azure Active Directory Managed Service Identity (MSI) gives your code an automatically managed identity for authenticating to Azure services, so that you can keep credentials out of your code.

Note: Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme.

Incorrect Answers:

A: Use the authentication-basic policy to authenticate with a backend service using Basic authentication. This policy effectively sets the HTTP Authorization header to the value corresponding to the credentials provided in the policy.

B: Anonymous is no authentication at all.

D: Your code needs credentials to authenticate to cloud services, but you want to limit the visibility of those credentials as much as possible. Ideally, they never appear on a developer's workstation or get checked-in to source control. Azure Key Vault can store credentials securely so they aren't in your code, but to retrieve them you need to authenticate to Azure Key Vault. To authenticate to Key Vault, you need a credential! A classic bootstrap problem.

Reference:

<https://azure.microsoft.com/en-us/blog/keep-credentials-out-of-code-introducing-azure-ad-managed-service-identity/>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-authentication-policies>


Community vote distribution

C (100%)

 **clarionprogrammer** Highly Voted 4 years, 2 months ago

C is correct.

upvoted 28 times


 **nnvuf** Highly Voted 4 years ago

A and D will send credential to the API.

B is not meet the requirement.

only C can be selected

upvoted 21 times

 **Vichu\_1607** Most Recent 8 months, 3 weeks ago

**Selected Answer: C**

Managed identities provide an identity for applications to use when connecting to resources that support Azure Active Directory (Azure AD) authentication. They are automatically managed by Azure and enable you to authenticate to services that support Azure AD authentication, without needing to insert credentials into your code.

In this case, a managed identity would meet both requirements: it would authenticate all API calls, and callers to the API would not need to send credentials because the authentication is handled automatically by Azure

upvoted 1 times

 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: C**

C. Managed identity

upvoted 2 times

🗨️ 👤 **gmishra88** 2 years, 8 months ago

Another pointless requirement made up with some answer in mind and it is impossible to guess what the Microsoft guy had in his mind at that point in time (of course nobody reviews relevance of the questions once created). The requirement that the credentials are not sent is weird. It depends on whether a client certificate is seen as a credential. That gets in conflict with whether the API is called from a client that does not have a managed identity. For example, from postman or curl or a Single Page Application. You cannot use Managed identity in that case. So, depending on how much an unfortunate question taker knows there is 50% chance (excuse my probability knowledge) that he can make a mistake or not. Microsoft guy wins

upvoted 5 times

🗨️ 👤 **gmishra88** 2 years, 8 months ago

My mistake, the requirement did say one azure resource to access another azure resource. In that case Managed Identity is clearly a winner. One must really read the question properly

upvoted 5 times

🗨️ 👤 **florianwicher** 4 months, 3 weeks ago

In my view, your original criticism is still valid. The scenario of "API accessing other resource" still leaves open the question of who the resource identifies as. One could imagine delegation, where the API acts on behalf of a client external to Azure; or one could imagine the Azure API acting as itself.

upvoted 1 times

🗨️ 👤 **Nokaido** 2 years, 9 months ago

**Selected Answer: C**

Tested it with a App Service running my C# code while using the Azure identity packages.

upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

**Selected Answer: C**

C is the correct answer.

upvoted 2 times

🗨️ 👤 **Eltooth** 3 years ago

<https://docs.microsoft.com/en-us/azure/api-management/api-management-policies#authentication-policies>

upvoted 1 times

🗨️ 👤 **AZ204Cert** 3 years, 2 months ago

Got this on 04/05/22 (selected Managed Identity)

upvoted 7 times

🗨️ 👤 **petitbilly** 3 years, 3 months ago

Got it in exam 03/22

upvoted 2 times

🗨️ 👤 **ning** 3 years, 10 months ago

Question is about from APIM to other azure resource, C is correct managed identity. Question is NOT about how end users access APIM

upvoted 6 times

🗨️ 👤 **francis6170** 4 years ago

got this in the exam :)

upvoted 5 times

🗨️ 👤 **businesskasper** 4 years, 1 month ago

Managed identity only works for internal api calls. The requirement is "all api calls must be authenticated". Certificate would meet that requirement?

upvoted 4 times

🗨️ 👤 **glam** 4 years, 1 month ago

correct

upvoted 3 times

🗨️ 👤 **MrZoom** 4 years, 3 months ago

D would work though, and it meets all the criteria (the private key of the certificate is never sent to the server during https handshake). If multiple answers are possible, I'd go for C and D, but if only one answer is possible, C is better for the reason stated in the answer explanation section.

upvoted 2 times

🗨️ 👤 **MrZoom** 4 years, 3 months ago

Also, using certificates for authentication isn't typical use with API's hosted in Azure (which \_is\_ part of the question).

upvoted 2 times

  **Spooky7** 4 years, 1 month ago

How ClientCertificate would help Azure API to authenticate in other Azure resources? The question is not about client authentication but Azure API authentication to use other Azure resources.

upvoted 4 times

## DRAG DROP -

You are developing an application. You have an Azure user account that has access to two subscriptions.

You need to retrieve a storage account key secret from Azure Key Vault.

In which order should you arrange the PowerShell commands to develop the solution? To answer, move all commands from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

**Powershell commands****Answer Area**

```
$secretvalue = ConvertTo-SecureString
$storAcctkey -AsPlainText
-Force
    Set-AzKeyVaultSecret -VaultName
$vaultName -Name $secretName
-SecretValue $secretvalue
```

```
Get-AzStorageAccountKey -
ResourceGroupName $resGroup -Name
$storAcct
```

```
Set-AzContext -SubscriptionId
$subscriptionID
```

```
Get-AzKeyVaultSecret -VaultName
$vaultName
```

```
Get-AzSubscription
```

**Suggested Answer:****Powershell commands****Answer Area**

```
$secretvalue = ConvertTo-SecureString
$storAcctkey -AsPlainText
-Force
    Set-AzKeyVaultSecret -VaultName
$vaultName -Name $secretName
-SecretValue $secretvalue
```

```
Get-AzStorageAccountKey -
ResourceGroupName $resGroup -Name
$storAcct
```

```
Set-AzContext -SubscriptionId
$subscriptionID
```

```
Get-AzKeyVaultSecret -VaultName
$vaultName
```

```
Get-AzSubscription
```

```
Get-AzSubscription
```

```
Set-AzContext -SubscriptionId
$subscriptionID
```

```
Get-AzStorageAccountKey -
ResourceGroupName $resGroup -Name
$storAcct
```

```
$secretvalue = ConvertTo-SecureString
$storAcctkey -AsPlainText
-Force
```

```
    Set-AzKeyVaultSecret -VaultName
$vaultName -Name $secretName
-SecretValue $secretvalue
```

```
Get-AzKeyVaultSecret -VaultName
$vaultName
```



Step 1: Get-AzSubscription -

If you have multiple subscriptions, you might have to specify the one that was used to create your key vault. Enter the following to see the subscriptions for your account:

Get-AzSubscription -

Step 2: Set-AzContext -SubscriptionId

To specify the subscription that's associated with the key vault you'll be logging, enter:

Set-AzContext -SubscriptionId <subscriptionID>

Step 3: Get-AzStorageAccountKey -

You must get that storage account key.

Step 4: \$secretvalue = ConvertTo-SecureString <storageAccountKey> -AsPlainText -Force

Set-AzKeyVaultSecret -VaultName <vaultName> -Name <secretName> -SecretValue \$secretvalue

After retrieving your secret (in this case, your storage account key), you must convert that key to a secure string, and then create a secret with that value in your key vault.




Step 5: Get-AzKeyVaultSecret -

Next, get the URI for the secret you created. You'll need this URI in a later step to call the key vault and retrieve your secret. Run the following PowerShell command and make note of the ID value, which is the secret's URI:

Get-AzKeyVaultSecret -VaultName <vaultName>

Reference:

<https://docs.microsoft.com/bs-latn-ba/Azure/key-vault/key-vault-key-rotation-log-monitoring>

  **AnonymousJhb**  4 years, 6 months ago

Answer looks correct

upvoted 60 times

  **Bear\_Polar** 2 years, 7 months ago

Correct! Reference: <https://learn.microsoft.com/bs-latn-ba/azure/key-vault/secrets/tutorial-rotation-dual?tabs=azurepowershell#add-the-storage-account-access-keys-to-key-vault-secrets>

upvoted 4 times

  **Frakandel**  4 years, 1 month ago

Question does not explain where the variables come from... Assuming I set the variables myself (as a "side effect" of the command), the solution below does make sense:



```
1. Get-AzSubscription
2. Set-AzContext -SubscriptionId $subscriptionID
3. Get-AzKeyVaultSecret -VaultName $vaultName
4. Get-AzStorageAccountKey -ResourceGroupName $resGroup -Name $storAcct
5. $secretvalue = ConvertTo-SecureString $storAcctkey -AsPlainText -Force
Set-AzKeyVaultSecret -VaultName $vaultName -Name $secretName -SecretValue $secretvalue
```

upvoted 46 times

  **zolept**  7 months, 1 week ago

correct

upvoted 1 times

  **p2006** 1 year, 9 months ago

<https://learn.microsoft.com/en-us/azure/key-vault/secrets/tutorial-rotation-dual?tabs=azurepowershell#add-storage-account-access-key-to-key-vault-secrets>

upvoted 4 times

  **suku555** 3 years, 1 month ago



Given Answer is correct only. "Get-AzKeyVaultSecret -VaultName \$vaultName" is just to make sure that we stored the key in vault (verification/double check). Placing this command anywhere else in the order does not make any sense.

upvoted 3 times

  **suku555** 3 years, 1 month ago

Also the main purpose of "Get-AzKeyVaultSecret -VaultName \$vaultName" at last is for the future usage..

upvoted 2 times

  **Azprep** 3 years, 2 months ago

Given answer us correct

upvoted 2 times

  **MiraA** 3 years, 9 months ago

I believe the assignment wants to select a proper subscription, then to retrieve the storage account key, then to store this key into the KeyVault and finally to check the secret was inserted properly. So:

1. Get-AzSubscription (\$subscriptionID = (...).SubscriptionName)
2. Set-AzContext
3. Get-AzStorageAccountKey (\$storAcctkey = "(...)[0].Value")
4. ConvertTo-SecureString, Set-AzKeyVaultSecret
5. Get-AzKeyVaultSecret (list the secrets stored in the KeyVault - check only)

upvoted 31 times

🗨️ 👤 **NPE\_** 1 year, 7 months ago

Thank you so much for the explanation!!! I was totally lost after seeing the answer.

upvoted 1 times

🗨️ 👤 **ning** 3 years, 10 months ago

Does the question miss some part? Why it has to save the key before retrieval? No where in the question mentioned that ... First two Get-AzSubscription

Set-AzContext - SubscriptionId \$subscriptionID

are absolutely correct, since there are two subscriptions, you have to point to the correct one ... but I am confused with 3, 4, 5 steps, what are those?

upvoted 13 times

🗨️ 👤 **glam** 4 years, 1 month ago

correct.

upvoted 3 times

🗨️ 👤 **sushikid112** 4 years, 5 months ago

Agree with some people saying the question and picture does not make sense. The question only says how to retrieve the secret from Key Vault. Based on the options, we also have to retrieve the value and insert it to Key Vault first. The question/picture lacks information based on what we have.

upvoted 9 times

🗨️ 👤 **Rodashar** 4 years, 6 months ago

Is this question actually on the exam? Even the url provided to explain the answer says this code is meant to cycle the storage keys stored in the vault not retrieve a key from the vault. Either the question is wrong or the answer is wrong.

upvoted 8 times

🗨️ 👤 **Tealon** 4 years, 6 months ago

I agree, the question does not match the options to give the answer. If the question was to store the account key to the key vault and then retrieve it from the key vault, the answer would be correct.

upvoted 6 times

🗨️ 👤 **iiiihhhh** 4 years, 6 months ago

There is another discussion for this question: <https://www.examttopics.com/discussions/microsoft/view/22272-exam-az-204-topic-3-question-5-discussion/>

upvoted 2 times

🗨️ 👤 **iiiihhhh** 4 years, 6 months ago

The question is: How to get secret from key vault, when there are two subscriptions available?

Available components(in order of boxes):

1. converting \$storageAcctkey to secure string and storing az secret with name \$secretName
2. getting storage account key from storage account (but without storing to variable \$storageAcctkey)
3. setting subscription context for use in current session (<https://docs.microsoft.com/en-us/powershell/module/az.accounts/set-azcontext?view=azps-5.1.0>)
4. list secrets in the vault
5. list subscriptions available

So for retrieving secret we (I think) need 5,3,4 and we need also SecureStringToBSTR (<https://docs.microsoft.com/cs-cz/azure/key-vault/secrets/quick-create-powershell>).

However the question seems to require using all actions available.

Is the question text really matching the picture?

upvoted 7 times

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 develop Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.

You need to obtain an Azure Resource Manager access token.

Solution: Use an X.509 certificate to authenticate the VM with Azure Resource Manager.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**




Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

Community vote distribution

B (100%)

  **KeerthiKP**  4 years, 9 months ago

No is the right answer:




Explanation :

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

link : <https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>  
upvoted 31 times

  **clarionprogrammer** 4 years, 2 months ago

Using Invoke-WebRequest vs Invoke-RestMethod doesn't matter. The point is that using a certificate is wrong for Identity Management.  
upvoted 13 times

  **mlantonis**  4 years, 1 month ago

Correct Answer: B - No




A certificate cannot be used to authenticate.

Instead run the Invoke-RestMethod or Invoke-WebRequest cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/overview>

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>  
upvoted 16 times

  **Vichu\_1607**  8 months, 3 weeks ago

**Selected Answer: B**

B. No

The solution does not meet the goal. While you can use an X.509 certificate to authenticate a service principal in Azure AD, this is not the recommended way to grant a VM access to specific resource groups in Azure Resource Manager.

Instead, you should use Managed Identities for Azure resources. A managed identity is an identity registered in Azure Active Directory that is



automatically managed by Azure. You can use this identity to authenticate to any service that supports Azure AD authentication, including Azure Resource Manager, without any credentials in your code.

upvoted 2 times

🗨️ 👤 **Razvan\_C** 1 year, 7 months ago

Is this still a valid question for the exam? I remember VMs are not in the study guide anymore.

upvoted 4 times

🗨️ 👤 **Esward** 2 years, 5 months ago

Correct Answer: B - No

Azure RBAC is an authorization system built on Azure Resource Manager that provides fine-grained access management of Azure resources. Here we need to authenticate.

Here we need to make use of managed identities for the virtual machine. Role-based access control is used for authorization and not authentication.

Instead run the `Invoke-RestMethod` or `Invoke-WebRequest` cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/overview>

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

upvoted 3 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

Correct Answer: B

upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

**Selected Answer: B**

B is the correct answer.

upvoted 1 times

🗨️ 👤 **glam** 4 years, 1 month ago

B. No...

upvoted 2 times

🗨️ 👤 **Frakandel** 4 years, 1 month ago

Yes, because a certificate can be used to authenticate... See also: <https://devblogs.microsoft.com/premier-developer/centralized-vm-certificate-deployment-across-multiple-regions-with-arm-templates/>

upvoted 2 times

🗨️ 👤 **TakumaK** 4 years, 1 month ago

Yes, your right. But can you explain this part in the question?

"You need to obtain an Azure Resource Manager access token"

upvoted 2 times

🗨️ 👤 **130nk3r5** 1 year, 6 months ago

Sure, I'd be happy to explain.

An Azure Resource Manager access token is a type of security token that is used to authenticate requests to the Azure Resource Manager API. This API is used to manage resources in Azure, such as virtual machines, storage accounts, and databases.

When you make a request to the Azure Resource Manager API, you need to include an access token in the Authorization header of your HTTP request. This token proves that you are who you say you are and that you have permission to perform the requested operation.

You can obtain an access token by authenticating with Azure Active Directory (Azure AD). There are several ways to do this, including using a client secret, a certificate, or a managed identity.



Once you have an access token, you can use it to authenticate requests to the Azure Resource Manager API. This allows you to manage resources in Azure programmatically, without needing to manually log in through the Azure portal.

upvoted 1 times

  **RahulKate** 4 years, 6 months ago

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

upvoted 1 times

  **Tealon** 4 years, 6 months ago

The given answer is correct.

upvoted 5 times

  **cyberbull** 4 years, 9 months ago

No , is the correct Answer

upvoted 8 times

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 develop Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.

You need to obtain an Azure Resource Manager access token.

Solution: Use the Reader role-based access control (RBAC) role to authenticate the VM with Azure Resource Manager.

Does the solution meet the goal?

A. Yes

B. No


**Suggested Answer: B**

Instead run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

Community vote distribution



**mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: B - No

Azure RBAC is an authorization system built on Azure Resource Manager that provides fine-grained access management of Azure resources. Here we need to authenticate.

Here we need to make use of managed identities for the virtual machine. Role-based access control is used for authorization and not authentication.

Instead run the Invoke-RestMethod or Invoke-WebRequest cmdlet to make a request to the local managed identity for Azure resources endpoint.

Reference:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/overview>

[https://docs.microsoft.com/en-us/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm](https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm)

upvoted 39 times

**overhill** 8 months ago

Phenomenal

upvoted 1 times

**KeerthiKP** Highly Voted 4 years, 9 months ago

ITs not a invoke-RestMethod , it is: Invoke-WebRequest cmdlet

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

upvoted 21 times

**vizay** 1 month, 3 weeks ago

Invoke-RestMethod

upvoted 1 times

**oversio** 4 years, 8 months ago

Yes! you all right

upvoted 2 times

**Tom87** 4 years, 2 months ago

Both Invoke-RestMethod and Invoke-WebRequest can be used. Their behaviour is very similar.

<https://blog.truesec.com/2020/07/29/invoke-webrequest-vs-invoke-restmethod/>

upvoted 5 times

🗄️ 👤 **clarionprogrammer** 4 years, 2 months ago

Agreed. The real reason it is wrong is RBAC can't be used for Identity Management.

upvoted 2 times

🗄️ 👤 **clarionprogrammer** 4 years, 2 months ago

It's true that the Reader role is required so the VM can read/access the Azure Resource Manager. But, it's not used to authenticate the VM with the Azure Resource Manager.

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

upvoted 3 times

🗄️ 👤 **hubekpeter** Most Recent 2 years, 7 months ago

Grant access to specific resource groups in Azure Resource Manager - resource based access control is perfectly usable, but a key here is to limit the access to a specific group, which is not possible through reader role, you need to implement ABAC policy in this case. So the answer is NO.

upvoted 1 times

🗄️ 👤 **coffecold** 2 years, 8 months ago

Selected Answer: A

Answer A - because this particular setup is appropriate to get the token (among others).

User-assigned managed identity with RBAC is just one way to obtain an Azure Resource Manager access token. Having done that still you have to get the access token by an Invoke-WebRequest.

step 1 : get the access token

```
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&client_id=<value of the user-assigned managed identity>&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
```

step 2 : access Azure Resource Manager with the access token

See <https://learn.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-ua-arm>

upvoted 1 times

🗄️ 👤 **coffecold** 2 years, 8 months ago

although ... the Reader role-based access control (RBAC) role is not used here to get the token, it's the client ID in the invoke web-request. So to be save : answer B

upvoted 1 times

🗄️ 👤 **Eltooth** 3 years ago

Selected Answer: B

B is correct answer.

upvoted 3 times

🗄️ 👤 **sawipf270** 3 years, 6 months ago

Selected Answer: B

correct

upvoted 1 times

🗄️ 👤 **glam** 4 years, 1 month ago

B. No...

upvoted 2 times

🗄️ 👤 **ccinetto** 4 years, 6 months ago

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

Here it says that you first need to go to the Resource Group, and in the IAM tab grant access to the VM as a Reader Role

Then, when you are inside the VM, you can use the command Invoke-WebRequest cmdlet in order to get the Token

4 is Correct because it is the first step you need to do in order to get the Token


and question 5 is correct too

upvoted 4 times

🗄️ 👤 **azahran** 4 years, 5 months ago

The question is you need an access token? Q 4 is preparation step and did not retrieve an access token.

upvoted 3 times

  **bugimachi** 4 years, 6 months ago

I'm not sure. What you do in IAM / RBAC is authorizing (!) the VM to access the resource group. You do NOT use IAM / RBAC for authentication.



I mean, you are right in saying that it is needed to grant access for the VM. But strictly speaking, the question asks for authentication only.

upvoted 4 times

  **Gathix444** 4 years, 5 months ago


Yes you do. <https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-linux-vm-access-arm>

upvoted 1 times

  **27close** 4 years, 7 months ago

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager. answer is NO

upvoted 2 times

  **blueturtle** 4 years, 7 months ago

Yup the updated version is to use Invoke-WebRequest cmdlet.

upvoted 1 times

## HOTSPOT -

You are building a website that is used to review restaurants. The website will use an Azure CDN to improve performance and add functionality to requests.

You build and deploy a mobile app for Apple iPhones. Whenever a user accesses the website from an iPhone, the user must be redirected to the app store.

You need to implement an Azure CDN rule that ensures that iPhone users are redirected to the app store.

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```

"conditions": [ {
  "name": "IsDevice",
  "parameters": {
    "@odata.type": "#Microsoft.Azure.Cdn.Models.",
    "operator": "Equal",
    "matchValues": [ "
  } },
  {
    "name": "RequestHeader",
    "parameters": {
      "@odata.type": "#Microsoft.Azure.Cdn.Models.",
      "operator": "Contains",
      "selector": "
    "matchValues": [ "
  } }
]

```

iOS  
 Mobile  
 iPhone  
 Desktop

DeliveryRuleDeviceConditionParameters  
 DeliveryRuleCookiesConditionParameters  
 DeliveryRulePostArgsConditionParameters  
 DeliveryRuleRequestHeaderConditionParameters

FROM  
 PRAGMA  
 X-POWERED-BY  
 HTTP\_USER\_AGENT

DeliveryRuleDeviceConditionParameters  
 DeliveryRuleCookiesConditionParameters  
 DeliveryRulePostArgsConditionParameters  
 DeliveryRuleRequestHeaderConditionParameters

iOS  
 Mobile  
 iPhone  
 Desktop

**Answer Area**

Box 1: iOS -

Android

➤ iOS

☞ Windows Phone

Windows

macOS

### Box 2: DeliveryRulesDeviceConditionParameters

The `DeliveryRulesDeviceCondition` defines the `IsDevice` condition for the delivery rule. `parameters` defines the parameters for the condition.

### Box 3: HTTP\_USER\_AGENT -

Incorrect Answers:

- ⇒ The Pragma HTTP/1.0 general header is an implementation-specific header that may have various effects along the request-response chain. It is used for backwards compatibility with HTTP/1.0 caches.

☞ "X-Powered-By" is a common non-standard HTTP response header (most headers prefixed with an 'X-' are non-standard).

#### Box 4: DeliveryRuleRequestHeaderConditionParameters

DeliveryRuleRequestHeaderCondition defines the RequestHeader condition for the delivery rule. parameters defines the parameters for the condition.

### Box 5: iOS -

The Require approved client app requirement only supports the iOS and Android for device platform condition.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/concept-conditional-access-conditions>

<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/concept-conditional-access-grant>

 **Zidimirite** Highly Voted 4 years, 2 months ago

First box: Mobile

<https://docs.microsoft.com/en-us/python/api/azure-mgmt-cdn/azure.mgmt.cdn.models.isdevicematchconditionparametersmatchvaluesitem>

Third box: iPhone. We are not looking for other Apple devices running iOS. It's the iPhone we're looking for.

upvoted 97 times

🗨️ 👤 **titombo** 4 years, 2 months ago  
You are right, also the HTTP\_USER\_AGENT can be iPhone, iPad, etc... There isn't iOS. Even because iPhone and iPad I believe both could have iOS running in it.  
upvoted 3 times

🗨️ 👤 **atomicicebreaker** 4 years, 2 months ago  
I agree, question is about CDN and answer explanation is all about Azure AD...  
upvoted 4 times

🗨️ 👤 **minaritochuck** 3 years, 3 months ago  
yes for example:  
HTTP\_USER\_AGENT=Mozilla/5.0 (iPhone; U; CPU like Mac OS X; en)  
upvoted 4 times

🗨️ 👤 **baroo1** 2 years, 4 months ago  
First box: Mobile as per documentation provided by this user.  
  
Yup! Confirmed.  
Third box: iPhone.  
HTTP\_USER\_AGENT for iPhones is like the one below. iOS does not exist.  
Mozilla/5.0 (iPhone; U; CPU like Mac OS X; en) AppleWebKit/420+ (KHTML, like Gecko) Version/3.0 Mobile/1A543a Safari/419.3  
upvoted 7 times

🗨️ 👤 **kondapaturi** Highly Voted 4 years ago  
1 DeliveryRuleDeviceConditionParameters - we are first checking for a device condition, hence we need to use the condition of DeliveryRuleDeviceConditionParameters  
2.Mobile - The devices can be either Desktop or Mobile. These are the two accepted values. Here since we need to route requests based on mobile devices, we need to choose the value of Mobile.  
3.DeliveryRequestHeaderConditionParameters. we need to understand the type of operating system running on the device. We can get this information from the request headers. Hence, we need to use the parameter of DeliveryRequestHeaderConditionParameters.  
4.HTTP\_USER\_AGENT - we can check the HTTP\_USER\_AGENT property in the request header. In the user agent property of the request header, you will normally get information about the environment where the request is originating from. An example is given below where I am showing the request header from my own machine when I browse to a site.  
5.iOS - we need to check the operating system which will be iOS.  
upvoted 43 times

🗨️ 👤 **Knightie** 2 years, 10 months ago  
Only this one is correct, the value only takes Desktop/Mobile, last one is iOS, Linux...etc  
upvoted 2 times

🗨️ 👤 **ivan0590** 3 years, 2 months ago  
I agree on the first four points, but the fifth point is incorrect.

The correct answer is iPhone.  
Here you can see lots of iOS User Agent examples:  
[https://developers.whatismybrowser.com/useragents/explore/operating\\_system\\_name/ios/](https://developers.whatismybrowser.com/useragents/explore/operating_system_name/ios/)  
"iPhone" appears all the time and "Mac OS X" is used instead of "iOS"  
  
Example:  
Mozilla/5.0 (iPhone; CPU iPhone OS 9\_3\_1 like Mac OS X) AppleWebKit/601.1.46 (KHTML, like Gecko) Version/9.0 Mobile/13E238 Safari/601.1  
upvoted 9 times

🗨️ 👤 **MiiExamT** Most Recent 3 months, 4 weeks ago  
For info:  
HTTP\_USER\_AGENT=Mozilla/5.0 (iPhone; U; CPU like Mac OS X; en) AppleWebKit/420+ (KHTML, like Gecko) Version/3.0 Mobile/1C25 Safari/419.3  
  
Last box should be iphone  
upvoted 1 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 6 months ago  
I got this question on my exam, 2023Dec, go with what I remember was the most voted answer. Score 902, most of the questions were here, slightly different on wording because the Azure Ad <-> Entra Id change. Case was City Power & Light. Good luck!  
Important tip, you have access to microsoft learn during the exam!



upvoted 4 times

🗨️ 👤 **giuliohome** 12 months ago

This quote is both misleading and dishonest: "You have access to Microsoft Learn during the exam." Such messages should be deleted by the admins, but there are many similar ones in the forum.

upvoted 2 times

🗨️ 👤 **Ahsanz10** 10 months, 1 week ago

but isn't it true, I believe we now have access to Microsoft Learn

upvoted 2 times

🗨️ 👤 **kotireddy4120** 1 year, 7 months ago

3rd box is iphone

upvoted 1 times

🗨️ 👤 **methkill** 1 year, 7 months ago

5th - the last one

upvoted 1 times

🗨️ 👤 **Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I used the examtopics answer, but the last one I choose iPhone - score 850

upvoted 1 times

🗨️ 👤 **Dixavado** 1 year, 9 months ago

It was on my exam today (2023-09-26) I used the examtopics answer - score 850

upvoted 1 times

🗨️ 👤 **nikipediaa** 1 year, 9 months ago

Question was on exam 2023-09-26

upvoted 1 times

🗨️ 👤 **Saluk\_DE** 2 years, 3 months ago

1. DeliveryRuleDeviceConditionParameters
2. Mobile
3. DeliveryRequestHeaderConditionParameters
4. HTTP\_USER\_AGENT
5. iPhone

Question was on exam 2023-03-30

upvoted 13 times

🗨️ 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Chosen DeviceRule, Mobile, Header, HTTP Request, iOS. Got 846. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 2 times

🗨️ 👤 **mabdo** 2 years, 4 months ago

was on exam 02/2023

upvoted 1 times

🗨️ 👤 **mabdo** 2 years, 4 months ago

on exam 02/23

upvoted 1 times

🗨️ 👤 **kociuba** 2 years, 6 months ago

got it on my exam 30-12-2022

score: 818

upvoted 1 times

🗨️ 👤 **vcfvct** 2 years, 8 months ago

Got this question on 10/21/2022

upvoted 3 times

🗨️ 👤 **CloudixExamTopics** 2 years, 7 months ago

Answer ?

upvoted 1 times

🗨️ 👤 **gmishra88** 2 years, 8 months ago

Really, I should know this to be a developer? This is so specific, Microsoft  
upvoted 10 times

🗨️ 👤 **DonH** 2 years ago

I agree this is insane.. how could they ask us to remember all this specific stuff.. I mean look at this table alone: <https://learn.microsoft.com/en-us/dotnet/api/microsoft.azure.management.cdn.models?view=azure-dotnet>  
and that's only a single question.. :(  
upvoted 1 times

🗨️ 👤 **Eltooth** 3 years ago

DeviceConditions  
Mobile  
HeaderConditions  
HTTP  
IOS  
upvoted 4 times

🗨️ 👤 **tosm** 3 years ago

reference: <https://docs.microsoft.com/en-us/azure/cdn/cdn-standard-rules-engine-match-conditions> the filter rule for device may only accept the value of mobile and desktop. So mobile/iphone is my vote.  
upvoted 1 times

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 are developing a website that will run as an Azure Web App. Users will authenticate by using their Azure Active Directory (Azure AD) credentials.

You plan to assign users one of the following permission levels for the website: admin, normal, and reader. A user's Azure AD group membership must be used to determine the permission level.

You need to configure authorization.

Solution:

- ⇒ Configure and use Integrated Windows Authentication in the website.
- ⇒ In the website, query Microsoft Graph API to load the groups to which the user is a member.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Microsoft Graph is a RESTful web API that enables you to access Microsoft Cloud service resources.


Instead in the Azure AD application's manifest, set value of the groupMembershipClaims option to All. In the website, use the value of the groups claim from the

JWT for the user to determine permissions.

Reference:

<https://blogs.msdn.microsoft.com/waws/2017/03/13/azure-app-service-authentication-aad-groups/>

Community vote distribution



**mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: B - No

Reference:

<https://docs.microsoft.com/en-us/archive/blogs/waws/azure-app-service-authentication-aad-groups>

<https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication>

<https://docs.microsoft.com/en-us/graph/api/resources/azure-ad-overview?view=graph-rest-1.0>

upvoted 36 times

**lugospod** 3 years, 5 months ago

I would say yes. Because if you have AD federation, then just by enabling WINDOWS authentication in your web app you gain access to Azure AD... so it is feasible

upvoted 1 times

**mattvasc** 3 years, 4 months ago

Authentication isn't the same as authorization, what about the roles? The answer is No.

upvoted 5 times

**Tom87** Highly Voted 4 years, 2 months ago

I would say the solution meets the goal, according to this:

<https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication>

upvoted 10 times

**coffecold** 2 years, 8 months ago

No, that is only for desktop or mobile application that runs on Windows. This is a web app in Azure.

upvoted 4 times

  **warchoon** 2 years, 4 months ago

The browser is such an application  
upvoted 1 times

  **Vichu\_1607** Most Recent 8 months, 2 weeks ago

**Selected Answer: B**

No, Solution not meet the goal  
upvoted 1 times

  **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023  
upvoted 5 times

  **hubekpeter** 2 years, 7 months ago

**Selected Answer: A**

It's perfectly doable. Voting for A.  
upvoted 1 times

  **vizay** 1 month, 3 weeks ago

While querying Microsoft Graph API to check the user's group membership is a valid approach, the first step – using Integrated Windows Authentication (IWA) – is not suitable for most Azure Web Apps that are accessed over the internet.

Here's why:

IWA is mainly for internal networks (like inside a company) – not for internet-based apps like Azure Web Apps.

Azure Web Apps typically use OpenID Connect/OAuth 2.0 with Azure AD for authentication – not IWA.

Also, IWA doesn't work cross-platform or on non-Windows devices easily.

✓ A better approach would be:



Register your web app in Azure AD.

Assign users to Azure AD groups (admin, normal, reader).

In your app, decode the JWT token and check the groups claim or call Microsoft Graph to get group membership.

Based on the group, set the user's permission.

upvoted 1 times

  **OPT\_001122** 2 years, 7 months ago

**Selected Answer: B**

Correct Answer: B. No  
upvoted 1 times

  **Eltooth** 3 years ago

**Selected Answer: B**

B is the correct answer.  
No.  
upvoted 2 times

  **vavra** 3 years ago

I think it's B, because of this:

<https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication>

Federated users only, i.e. those created in an Active Directory and backed by Azure Active Directory. Users created directly in AAD, without AD backing - managed users - cannot use this auth flow. This limitation does not affect the Username/Password flow.

upvoted 2 times

  **minaritochuck** 3 years, 1 month ago



Integrated Windows authentication is for "Intranet Application" (Although AD manages it) so don't use it on website  
upvoted 2 times

  **minaritochuck** 3 years, 3 months ago

It makes sense for app to use Integrated Windows Authentication (Implicit grant) for authentication and authorized by Microsoft Graph to access AD resources.

<https://docs.microsoft.com/en-us/graph/api/resources/directoryrole?view=graph-rest-1.0>

upvoted 1 times

  **EwanoE** 3 years, 3 months ago

**Selected Answer: B**

Question says "website that will run as an Azure Web App" and there's no Integrated Windows Authentication in Azure Web App (aka App Service)

upvoted 2 times

  **Freidrich** 3 years, 4 months ago

**Selected Answer: B**

The correct answer is B: No.

upvoted 1 times

  **mattdober** 3 years, 4 months ago

**Selected Answer: B**

I think its B

upvoted 1 times

  **ScubaDiver123456** 3 years, 5 months ago

**Selected Answer: A**

The answer is YES, assuming the application is hosted on a domain-joined VM. This document shows you how to set it up

<https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication>

upvoted 1 times

  **ScubaDiver123456** 3 years, 5 months ago

Sorry. It does say it is an Azure Web App which can only run on Azure App Service. You can't join the underlying VMs to Active Directory. Therefore, I'm changing my vote to NO

upvoted 9 times

  **ning** 3 years, 10 months ago

No, windows integrated authentication is NOT azure AD, you need app registration to get Azure AD

upvoted 7 times

  **kondapaturi** 4 years ago

Answer – No, Here you need to create an application in Azure AD. Then set the groupMembershipClaims claims. Then inspect the token in the application to see if the user is part of that group.

upvoted 5 times

  **Idkhow** 4 years ago

seriously what's the real answer??

in my opinion i think it's a yesn't

Configure and use Integrated Windows Authentication in the website.

(what does it have to do with the Integrated Windows Authentication)

In the website, query Microsoft Graph API to load the group to which the user is a member.

(like what the explanation said that it works)

upvoted 2 times

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 develop Azure solutions.

You must grant a virtual machine (VM) access to specific resource groups in Azure Resource Manager.

You need to obtain an Azure Resource Manager access token.

Solution: Run the Invoke-RestMethod cmdlet to make a request to the local managed identity for Azure resources endpoint.

Does the solution meet the goal?

A. Yes

B. No

#### Suggested Answer: A

Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager

You will need to use PowerShell in this portion.

1. In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.
2. Enter in your Username and Password for which you added when you created the Windows VM.
3. Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.
4. Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.


Example:

```
$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}
```

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>

Community vote distribution



**mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: Yes

Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.

I believe we can you Invoke-RestMethod too.

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>  
upvoted 32 times

**princesskay** Highly Voted 4 years, 10 months ago

This does make sense since a web-request can be a rest call!

In the link below, it explains how web-request supports HTML content while Rest-Method supports JSON and XML content. The link provided for the solution is correct as well.

<https://superuser.com/questions/1235349/what-is-the-difference-between-invoke-webrequest-and-invoke-restmethod>  
upvoted 25 times

**Vichu\_1607** Most Recent 8 months, 2 weeks ago

Selected Answer: A

Yes, the solution does meet the goal.

The Invoke-RestMethod cmdlet can be used to make a request to the local managed identity for Azure resources endpoint (<http://169.254.169.254/metadata/identity/oauth2/token>). This endpoint can provide an Azure Resource Manager access token when called from an Azure VM with a managed identity enabled. This token can then be used to authenticate to the Azure Resource Manager API.

upvoted 3 times

**overhill** 8 months ago

Strong as usual  
upvoted 1 times

🗨️ 👤 **carlosghosn** 2 years, 7 months ago  
Got this in the exam today ! Nov 25, 2022  
upvoted 2 times

🗨️ 👤 **OPT\_001122** 2 years, 7 months ago  
**Selected Answer: A**  
Answer: Yes  
upvoted 1 times

🗨️ 👤 **Mousavi** 2 years, 9 months ago  
**Selected Answer: A**  
Absolutely the corrcr answer is A  
upvoted 2 times

🗨️ 👤 **Perplex** 2 years, 11 months ago  
**Selected Answer: A**  
Answer A - The 2 most upvoted comments say this is correct and they are right.  
upvoted 3 times

🗨️ 👤 **Azprep** 3 years, 2 months ago  
Answer is Yes  
upvoted 2 times

🗨️ 👤 **SivajiTheBoss** 3 years, 3 months ago  
**Selected Answer: B**  
Need to use Invoke-WebRequest cmdlet not Invoke-RestMethod  
Correct Answer: B  
Get an access token using the VM's system-assigned managed identity and use it to call Azure Resource Manager  
You will need to use PowerShell in this portion.  
1. In the portal, navigate to Virtual Machines and go to your Windows virtual machine and in the Overview, click Connect.  
2. Enter in your Username and Password for which you added when you created the Windows VM.  
3. Now that you have created a Remote Desktop Connection with the virtual machine, open PowerShell in the remote session.  
4. Using the Invoke-WebRequest cmdlet, make a request to the local managed identity for Azure resources endpoint to get an access token for Azure Resource Manager.  
Example:  
`$response = Invoke-WebRequest -Uri 'http://169.254.169.254/metadata/identity/oauth2/token?api-version=2018-02-01&resource=https://management.azure.com/' -Method GET -Headers @{Metadata="true"}`  
Reference:  
<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>  
upvoted 1 times

🗨️ 👤 **mabusalma** 3 years, 5 months ago  
**Selected Answer: A**  
The answer is Yes correct  
Invoke-WebRequest or Invoke-RestMethod will do the same job but the response in RestMethod will be much simpler.  
upvoted 3 times

🗨️ 👤 **LauraGF** 3 years, 10 months ago  
The answer is no  
Invoke-RestMethod: Grant your VM access to a secret stored in a Key Vault  
<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-nonaad>  
  
Invoke-WebRequest: Grant your VM access to a Resource Group in Azure Resource Manager  
<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/tutorial-windows-vm-access-arm>  
upvoted 3 times

🗨️ 👤 **kondapaturi** 4 years ago  
Answer is YES, you can make use of the PowerShell commands to make a request to managed identity service of the virtual machine. The Invoke-RestMethod is a generic call that is available in PowerShell for making a call to a web service.

upvoted 4 times

🗨️ 👤 **Spooky7** 4 years, 1 month ago

In order to grant VM access to ARM you need to:

- assign VM user/system managed identity
- configure permission to ARM for that identity
- and then you can access token by Invoke-WebRequest cmdlet and use to to authenticate in ARM.

There is nothing in the question which indicates that those 2 previous steps are already done, so I would say it is not a correct answer.

upvoted 4 times

🗨️ 👤 **glam** 4 years, 1 month ago

A. Yes

upvoted 2 times

🗨️ 👤 **clarionprogrammer** 4 years, 2 months ago

A. Yes

Invoke-RestMethod is merely a rapper commandlet around Invoke-WebRequest.

upvoted 2 times

🗨️ 👤 **ferut** 4 years, 1 month ago

roflol... do you mean wrapper?

upvoted 2 times

🗨️ 👤 **businesskasper** 4 years, 1 month ago

I'm pretty sure its a rapper cmdlet

upvoted 1 times

🗨️ 👤 **Santileo** 4 years, 2 months ago

I code in powershell regularly and Invoke-RestMethod is the correct one to fetch these tokens.

upvoted 4 times

🗨️ 👤 **kwaazaar** 4 years, 2 months ago

Where does it say the code is running on the VM itself?

upvoted 1 times



## HOTSPOT -

You are building a website to access project data related to teams within your organization. The website does not allow anonymous access. Authentication is performed using an Azure Active Directory (Azure AD) app named internal.

The website has the following authentication requirements:

- ⇒ Azure AD users must be able to login to the website.
- ⇒ Personalization of the website must be based on membership in Active Directory groups.

You need to configure the application's manifest to meet the authentication requirements.

How should you configure the manifest? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
{
  ...
  "appId": "d61126e3-089b-4adb-b721-d5023213df7d",
  "displayName": "internal",
  [Box 1] : "All",
  "optionalClaims"
  "groupMembershipClaims"
  [Box 2] : true
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
  ...
}
```

### Answer Area

Suggested Answer:

```
{
  ...
  "appId": "d61126e3-089b-4adb-b721-d5023213df7d",
  "displayName": "internal",
  [Box 1] : "All",
  "optionalClaims"
  "groupMembershipClaims"
  [Box 2] : true
  "allowPublicClient"
  "oauth2Permissions"
  "requiredResourceAccess"
  "oauth2AllowImplicitFlow"
  ...
}
```

Box 1: groupMembershipClaims -

Scenario: Personalization of the website must be based on membership in Active Directory groups.

Group claims can also be configured in the Optional Claims section of the Application Manifest.

Enable group membership claims by changing the groupMembershipClaim

The valid values are:

- "All"
- "SecurityGroup"
- "DistributionList"
- "DirectoryRole"

Box 2: oauth2Permissions -

Scenario: Azure AD users must be able to login to the website. oauth2Permissions specifies the collection of OAuth 2.0 permission scopes that the web API (resource) app exposes to client apps. These permission scopes may be granted to client apps during consent.

Incorrect Answers:

oauth2AllowImplicitFlow. oauth2AllowImplicitFlow specifies whether this web app can request OAuth2.0 implicit flow access tokens. The default is false. This flag is used for browser-based apps, like Javascript single-page apps.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims>

🗨️ 👤 **AakashNeedsEmAll** Highly Voted 4 years, 7 months ago

Second answer is incorrect. Well, `oauth2Permissions` can only accept collections value like an array not a boolean. It should be `oauth2AllowImplicitFlow`.

upvoted 176 times

🗨️ 👤 **cbn** 4 years, 4 months ago

Agree.

`oauth2permissions`:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2permissions-attribute>

`oauth2allowimplicitflow`:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2allowimplicitflow-attribute>

upvoted 14 times

🗨️ 👤 **Asteriz** 3 years, 5 months ago

Agreed

upvoted 1 times

🗨️ 👤 **Juanlu** 4 years, 4 months ago

I Agree with you:

1) `groupMembershipClaims`

2) `oauth2AllowImplicitFlow`

upvoted 22 times

🗨️ 👤 **Esward** 2 years, 5 months ago

Agreed

upvoted 1 times

🗨️ 👤 **cloud\_exam1** 4 years, 7 months ago

I think so. `oauth2Permissions` is a list. Should be `oauth2AllowImplicitFlow`.

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest>

upvoted 11 times

🗨️ 👤 **iiiihhh** 4 years, 6 months ago

The problem is that `oauth2AllowImplicitFlow` is for browser-based apps(SPA), what is not mentioned in the question.

upvoted 1 times

🗨️ 👤 **Tealon** 4 years, 6 months ago

OAuth2 can also be used in this scenario. It does not matter if you use the normal grant or implicit flow.

upvoted 4 times

🗨️ 👤 **AfroYeti** 4 years, 5 months ago

It is "You are building a WEBSITE to access project data related to teams within your organization."

I would imagine that, that website would run in a browser

upvoted 9 times

🗨️ 👤 **mlantonis** Highly Voted 4 years, 1 month ago

Box 1: `groupMembershipClaims`

Personalization of the website must be based on membership in Active Directory groups.

Group claims can also be configured in the Optional Claims section of the Application Manifest. Enable group membership claims by changing the `groupMembershipClaim`

The valid values are:

- "All"

- "SecurityGroup"

- "DistributionList"

- "DirectoryRole"

Here we need to mention that we want to get the groups for the users. Hence we need to mention to set the `groupMembershipClaims` property to All.

Box 2: `oauth2AllowImplicitFlow`

Azure AD users must be able to login to the website.

auth2Permissions can only accept collections value like an array, not a boolean. oauth2AllowImplicitFlow accepts boolean value.

Here from the list of options given, if we want the application to fetch the required tokens , we would need to allow Implicit Flow.

upvoted 44 times

  **mlantonis** 4 years, 1 month ago

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#groupmembershipclaims-attribute>



<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2permissions-attribute>

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest#oauth2allowimplicitflow-attribute>

<https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-group-claims>

<https://docs.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-implicit-grant-flow>

upvoted 8 times

  **florianwicher** Most Recent 4 months, 3 weeks ago

Why not "allowPublicClient"? If some logic is implemented in Javascript and running in the browser, the app would count as a public client, and we'd potentially have to turn this flag on. It is certainly not a weirder choice than turning on the deprecated implicit flow.

upvoted 1 times

  **kotireddy4120** 1 year, 7 months ago

I Agree with :

1) groupMembershipClaims

2) oauth2AllowImplicitFlow

upvoted 3 times

  **Karunan** 1 year, 7 months ago

If its implicitflow, why answer not yet changed?

upvoted 2 times

  **Occ50bf** 10 months, 3 weeks ago

Because if this website had all the right answers, Microsoft would sue them.

upvoted 1 times

  **richardc7** 1 year, 9 months ago

Got this question on 28/09/2023

upvoted 1 times

  **BaoNguyen2411** 1 year, 11 months ago

got this question on 06-29-2023

upvoted 3 times

  **MysticalSam** 2 years ago

This question was in today's exam on 10-June-2023

upvoted 4 times

  **adilkhan** 2 years, 2 months ago

ChatGPT

1) groupMembershipClaims

2) oauth2AllowImplicitFlow

```
{
  "signInAudience": "AzureADMyOrg",
  "groupMembershipClaims": "SecurityGroup",
  "appId": "<app ID>",
  "appRoles": [],
  "availableToOtherTenants": false,
  "displayName": "<app name>",
  "oauth2AllowImplicitFlow": false,
  "optionalClaims": {
```

```
"idToken": [],
"accessToken": []
},
"publisherDomain": "<your domain>",
"requiredResourceAccess": []
}
```

upvoted 2 times

🗉 👤 **carlosghosn** 2 years, 7 months ago

Got this in the exam today ! Nov 25, 2022

upvoted 1 times

🗉 👤 **DusanS** 2 years, 8 months ago

It must be oAuth2AllowImplicitFlow because it is boll type in contrast to collection type. Look carefully picture.

upvoted 1 times

🗉 👤 **Eltooth** 3 years ago

1) groupMembershipClaims

2) oAuth2AllowImplicitFlow

upvoted 4 times

🗉 👤 **Evo\_Morales** 3 years, 2 months ago

Check Point - how many folks knew the correct answer before reading/researching?

upvoted 1 times

🗉 👤 **ReniRechner** 3 years, 3 months ago

According to

<https://docs.microsoft.com/en-us/azure/active-directory/develop/reference-app-manifest>

optionalClaims => "All" might not be allowed

groupMembershipClaims => "None", ... "All"

allowPublicClient => boolean

oAuth2Permissions => collection

requiredResourceAccess => collection

oAuth2AllowImplicitFlow => boolean

So first has to be "groupMembershipClaims" (can also be inferred from the requirements).

"allowPublicClient" seems to be meant for special cases not required here, there is only "oAuth2AllowImplicitFlow" which makes sense for a website.

upvoted 3 times

🗉 👤 **oescm** 3 years, 4 months ago

Got this one 02/2022. Went with:

1) groupMembershipClaims

2) oAuth2AllowImplicitFlow

upvoted 3 times

🗉 👤 **MasterQuestMaster** 3 years, 5 months ago

Got this on the exam!

upvoted 3 times

🗉 👤 **Iugospod** 3 years, 5 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 5 times

🗉 👤 **maharaju** 3 years, 2 months ago

How come after writing exams you guys are visiting each questions and commenting ..SO much time

upvoted 8 times

You develop an app that allows users to upload photos and videos to Azure storage. The app uses a storage REST API call to upload the media to a blob storage account named Account1. You have blob storage containers named Container1 and Container2.

Uploading of videos occurs on an irregular basis.

You need to copy specific blobs from Container1 to Container2 when a new video is uploaded.

What should you do?

- A. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API
- B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet
- C. Use AzCopy with the Snapshot switch to copy blobs to Container2
- D. Download the blob to a virtual machine and then upload the blob to Container2

**Suggested Answer: B**

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob.

Example 1: Copy a named blob -

```
C:\PS>Start-AzureStorageBlobCopy -SrcBlob "ContosoPlanning2015" -DestContainer "ContosoArchives" -SrcContainer "ContosoUploads"
```

This command starts the copy operation of the blob named ContosoPlanning2015 from the container named ContosoUploads to the container named


ContosoArchives.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0>

Community vote distribution

B (100%)

 **AnonymousJhb** Highly Voted 4 years, 6 months ago

B looks correct.

Common Blob storage event scenarios include image or video processing, search indexing, or any file-oriented workflow. Asynchronous file uploads are a great fit for events. When changes are infrequent, but your scenario requires immediate responsiveness, event-based architecture can be especially efficient.


<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

upvoted 37 times

 **Esward** 2 years, 5 months ago

I agree with you! B is correct answer

upvoted 2 times

 **mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: B


The Start-AzureStorageBlobCopy cmdlet starts to copy a blob.

Reference:

<https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0>

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

upvoted 16 times

 **Vichu\_1607** Most Recent 8 months, 2 weeks ago

**Selected Answer: B**

The correct answer is B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet.

Azure Event Grid can react to relevant events across both Azure and non-Azure services in near-real time fashion. You can create an Event Grid subscription on the blob storage account (Account1) to trigger a function when a new blob (video) is uploaded. This function can use the Start-AzureStorageBlobCopy cmdlet to start an asynchronous copy of the blob from Container1 to Container2.

Option A, using the Put Blob operation, would require downloading and re-uploading the blob, which is not efficient for large blobs like videos.

Option C, using AzCopy with the Snapshot switch, would create a read-only snapshot of the blob, not a writable copy.

Option D, downloading the blob to a virtual machine and then uploading it to Container2, would also be inefficient for large blobs and would require unnecessary compute resources.

upvoted 1 times

🗳️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 5 times

🗳️ 👤 **Saluk\_DE** 2 years, 3 months ago

**Selected Answer: B**

Question was on exam 2023-03-30

upvoted 1 times

🗳️ 👤 **NombreFalso** 2 years, 4 months ago

**Selected Answer: B**

Only B is automated, since this happens on an irregular basis it needs to be automated

upvoted 3 times

🗳️ 👤 **OPT\_001122** 2 years, 6 months ago

**Selected Answer: B**

B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet

upvoted 1 times

🗳️ 👤 **coffecold** 2 years, 8 months ago

**Selected Answer: B**

I think the handler then should be a powershell function ?

upvoted 2 times

🗳️ 👤 **gmishra88** 2 years, 8 months ago

I do get it is the correct answer but the option looked very weird when I first read it: "B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet". Probably Microsoft guy wanted that effect. I would have said event-grid topic with an Start-AzureStorageBlobCopy event handler. Microsoft wins again.

upvoted 6 times

🗳️ 👤 **Eltooth** 3 years ago

**Selected Answer: B**

B is the correct answer.

upvoted 3 times

🗳️ 👤 **arindam1989** 3 years, 2 months ago

**Selected Answer: B**

The automation of the process comes into play. Answer should be (B)

upvoted 1 times

🗳️ 👤 **AZ204Cert** 3 years, 2 months ago

Got this on 04/05/22 (selected B)

upvoted 2 times

🗳️ 👤 **Mev4953** 3 years, 5 months ago

**Selected Answer: B**

Correct Answer B

upvoted 2 times

🗳️ 👤 **ehurfheiz** 3 years, 5 months ago

**Selected Answer: B**

B looks correct



upvoted 1 times

🗳️ 👤 **hems4all** 4 years ago

The Start-AzureStorageBlobCopy cmdlet starts to copy a blob

The Put Blob operation creates a new block, page, or append blob, or updates the content of an existing block blob. So in the given options, only Start-AzureStorageBlobCopy cmdlet copies the blob

upvoted 3 times

  **glam** 4 years, 1 month ago



B. Create an Event Grid topic that uses the Start-AzureStorageBlobCopy cmdlet

upvoted 4 times

  **faizalzain** 4 years, 2 months ago

the answer is C in udemy

upvoted 3 times

  **AOE** 3 years, 10 months ago

You dont have snapshots in the question so C is definitely wrong.

upvoted 3 times

You are developing an ASP.NET Core website that uses Azure FrontDoor. The website is used to build custom weather data sets for researchers. Data sets are downloaded by users as Comma Separated Value (CSV) files. The data is refreshed every 10 hours. Specific files must be purged from the FrontDoor cache based upon Response Header values. You need to purge individual assets from the Front Door cache. Which type of cache purge should you use?

- A. single path
- B. wildcard
- C. root domain

**Suggested Answer: A**


These formats are supported in the lists of paths to purge:

- ⇒ Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /  
[1]
- ⇒ Wildcard purge: Asterisk (\*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /\* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /\*, for example, /pictures/\*.
- ⇒ Root domain purge: Purge the root of the endpoint with "/" in the path.

Reference:

<https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>

Community vote distribution



**mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: A

These formats are supported in the lists of paths to purge:

- Single path purge: Purge individual assets by specifying the full path of the asset (without the protocol and domain), with the file extension, for example, /pictures/strasbourg.png;
- Wildcard purge: Asterisk (\*) may be used as a wildcard. Purge all folders, subfolders, and files under an endpoint with /\* in the path or purge all subfolders and files under a specific folder by specifying the folder followed by /\*, for example, /pictures/\*.
- Root domain purge: Purge the root of the endpoint with "/" in the path.

Here since we just need to purge individual assets, we can just make use of the Single path purge as the format for purging assets.

Reference:

<https://docs.microsoft.com/en-us/azure/frontdoor/front-door-caching>

upvoted 32 times

**Tealon** Highly Voted 4 years, 6 months ago

The given answer is correct.

upvoted 26 times

**Vichu\_1607** Most Recent 8 months, 2 weeks ago

**Selected Answer: A**

The correct answer is A. Single path.

In Azure Front Door, you can purge cached content by specifying individual paths (URLs) to the content. This is known as a single path purge. This would be the most appropriate option if you need to purge specific files based on Response Header values.



Option B, wildcard, allows you to purge multiple paths that match a certain pattern, and option C, root domain, purges all cached content for a specific domain. These options would not be as precise as a single path purge for purging specific files.

upvoted 2 times

🗳️ 👤 **rahfa** 11 months, 2 weeks ago

**Selected Answer: B**

what about "Specific files must be purged from the FrontDoor cache based upon Response Header values". shouldn't be wildcard?

upvoted 1 times

🗳️ 👤 **Esward** 2 years, 5 months ago

Given answer is correct!

<https://learn.microsoft.com/en-us/azure/frontdoor/standard-premium/how-to-cache-purge>

upvoted 1 times

🗳️ 👤 **Lucifer14** 2 years, 6 months ago

In 28-12-2022 exam

upvoted 4 times

🗳️ 👤 **capable** 2 years, 7 months ago

GOT THIS QUE IN EXAM 27/11/2022

upvoted 2 times

🗳️ 👤 **TheExamMaster2020** 2 years, 7 months ago

Did my exam on 15th November 2022. This question was on it.

upvoted 1 times

🗳️ 👤 **Eltooth** 3 years ago

**Selected Answer: A**

A is correct answer.

Specific file = Single path

upvoted 4 times

🗳️ 👤 **[Removed]** 3 years, 2 months ago

**Selected Answer: A**

Got this on 04/29/2022. Correct Answer: A

upvoted 3 times

🗳️ 👤 **AZ204Cert** 3 years, 2 months ago

Got this on 04/05/22 (selected Single path)

upvoted 2 times

🗳️ 👤 **petitbilly** 3 years, 3 months ago

Got it in exam 03/22

upvoted 2 times

🗳️ 👤 **Manivannan19** 3 years, 4 months ago

Is Azure FrontDoor part of this exam?

upvoted 3 times

🗳️ 👤 **coffecold** 2 years, 8 months ago

CDN is , Azure FrontDoor includes CDN. CDN has the caching thing.

upvoted 1 times

🗳️ 👤 **Freidrich** 3 years, 4 months ago

**Selected Answer: A**

The proposed answer is correct A: Single path.

upvoted 1 times

🗳️ 👤 **ehurfheiz** 3 years, 5 months ago

**Selected Answer: A**



A is the correct answer

upvoted 2 times

🗳️ 👤 **kondapaturi** 4 years ago

Single Path is correct , we just need to purge individual assets , we can just make use of the Single path purge as the format for purging assets.

upvoted 1 times

  **UnknowMan** 4 years, 1 month ago

The given answer is correct.

upvoted 1 times

Your company is developing an Azure API.

You need to implement authentication for the Azure API. You have the following requirements:

All API calls must be secure.

⇒ Callers to the API must not send credentials to the API.

Which authentication mechanism should you use?


- A. Basic
- B. Anonymous
- C. Managed identity
- D. Client certificate

**Suggested Answer: C**

Use the authentication-managed-identity policy to authenticate with a backend service using the managed identity of the API Management service. This policy essentially uses the managed identity to obtain an access token from Azure Active Directory for accessing the specified resource. After successfully obtaining the token, the policy will set the value of the token in the Authorization header using the Bearer scheme. Reference:

<https://docs.microsoft.com/bs-cyrl-ba/azure/api-management/api-management-authentication-policies>

Community vote distribution



🗳️ 👤 **MasDen** Highly Voted 4 years, 7 months ago

If a caller outside Azure doesn't look like we can use Managed Identity in such a situation. I think the client certificate is better.  
upvoted 39 times

🗳️ 👤 **1CY1** 11 months, 1 week ago

Seems question is very unclear. You might be right but it also says not sending credentials so therefor C Managed Identity. It says '...must not send credentials to the API.', Client certificate is sending credentials to the API.  
upvoted 2 times

🗳️ 👤 **dluk** 4 years, 7 months ago

Seems it's legit because question says "All API calls must be secure"  
upvoted 2 times

🗳️ 👤 **IsildursHeir** 4 years, 6 months ago

They do say "Azure API" so I am guessing they mean its an Azure provided API, not one wrapped around by Azure API management so MI would work..  
upvoted 3 times

🗳️ 👤 **jay158** 4 years ago

Outsides user's can authenticate using Facebook /Google Identity providers.  
Managed Identity is a good answer.  
upvoted 5 times

🗳️ 👤 **ning** 3 years, 10 months ago

Those are NOT managed identity, those are claim based authentication.  
upvoted 7 times

🗳️ 👤 **Knightie** 2 years, 10 months ago

there is another similar question, using the cert, you will need to pass the cert or else store the cert in the key vault that you will still need to pass a value to retrieve the secret from key vault, so you will anyway pass a value to get it done. Managed ID doesn't need to pass a value like this.  
upvoted 1 times

🗳️ 👤 **luppittegui** Highly Voted 4 years, 6 months ago

Ans is correct: managed identities eliminate the need for developers having to manage credentials by providing an identity for the Azure resource in Azure AD and using it to obtain Azure Active Directory (Azure AD) tokens.

<https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>

upvoted 31 times

🗄️ 👤 **Vady98** 3 years, 8 months ago

If a caller outside Azure doesn't look like we can use Managed Identity.  
I think the client certificate is better.

upvoted 2 times

🗄️ 👤 **Vichu\_1607** Most Recent 8 months, 2 weeks ago

Selected Answer: C

The correct answer is C. Managed identity.

Managed identities for Azure resources is a feature of Azure Active Directory. It provides Azure services with an automatically managed identity in Azure AD. You can use this identity to authenticate to any service that supports Azure AD authentication, without having any credentials in your code.

Option A, Basic authentication, would require sending credentials (username and password) with each API call.

Option B, Anonymous authentication, would not secure the API calls as it allows anyone to call the API without any form of authentication.

Option D, Client certificate authentication, would require the callers to present a valid client certificate, which can be considered as a form of credential.

upvoted 1 times

🗄️ 👤 **sarmaria** 2 years, 3 months ago

Got this on 16/03/23. Went with proposed solution. Make sure to prepare for case study. I got city and lights case study. No Kubernetes, Search, Logic Apps questions for me.

upvoted 3 times

🗄️ 👤 **YLM007** 2 years, 1 month ago

have u gone through all the questions from exam topics?

upvoted 2 times

🗄️ 👤 **NombreFalso** 2 years, 4 months ago

Selected Answer: C

C because it doesn't send credentials + API that's gonna call it is on Azure so Managed Identity is viable.

upvoted 2 times

🗄️ 👤 **Esward** 2 years, 5 months ago

Given answer is correct as per MS docs

<https://learn.microsoft.com/en-us/azure/api-management/authentication-managed-identity-policy>

upvoted 1 times

🗄️ 👤 **gmishra88** 2 years, 8 months ago

It can never be managed identity unless the client is an azure resource that supports managed identity. The question does not say that. So, it has to be client-certificate. But then this is Microsoft exam and cannot expect a lot of logic from the question creators.

upvoted 2 times

🗄️ 👤 **bamosk** 2 years, 11 months ago

Selected Answer: C

C is correct. A lot of people choose D (Client Certificate), but that is also sent to the server using the header X-ARR-ClientCert so that is not correct.

See here: <https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth#access-client-certificate>

upvoted 2 times

🗄️ 👤 **Eltooth** 3 years ago

C appears best answer based on question text, however if question is missing key words (API Management) then D would be best answer.

Remember the exam question will not be ambiguous nor prove to someone's photo memory of rewriting it here.

upvoted 3 times

🗄️ 👤 **tosm** 3 years ago

my read was: my company develop an Azure API, and need to authenticate the API to other (Azure) services. If this is the case. using managed ID is correct.

upvoted 3 times

🗨️ **suku555** 3 years, 1 month ago

key words - developing an Azure API (obviously customization and there is no Azure service named 'Azure API') & Secure API calls --> These leads to the answer 'Client Certificate'.. Also in the Q there is no requirement of managing/rotating credential so surprising why many people vote for 'Managed Identity'..

upvoted 1 times

🗨️ **Azprep** 3 years, 2 months ago

Managed identity

upvoted 1 times

🗨️ **iamstudying** 3 years, 3 months ago

**Selected Answer: D**

Agree with MasDen

upvoted 2 times

🗨️ **petitbilly** 3 years, 3 months ago

Got it in exam 03/22

upvoted 1 times

🗨️ **DonOnur** 3 years, 4 months ago

**Selected Answer: C**

Since there is no requirement for users outside of Azure, I would choose C.

upvoted 4 times

🗨️ **DonOnur** 3 years, 4 months ago

Couldn't find the requirement that it has to be available to users outside azure.

upvoted 4 times

🗨️ **ytingyeu** 3 years, 4 months ago

There is no service called "Azure API". Thus, I assume it refers to "App Service > API Apps".

<https://azure.microsoft.com/en-us/services/app-service/api>

In the page of "Secure app > Use managed identities", it's all about App itself accesses other resources, not about client calling API.

<https://docs.microsoft.com/en-us/azure/app-service/overview-managed-identity>

And the page "AuthN and AuthZ in Azure App Service ", it says "App Service uses federated identity", not managed identities.

<https://docs.microsoft.com/en-us/azure/app-service/overview-authentication-authorization>

So I suppose the only thing we can do is client certificate

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth>

upvoted 3 times

🗨️ **ytingyeu** 3 years, 4 months ago

Plus, in page "Configure your App Service app to use AAD login" which is a part of federated identity, I don't see managed identities either.

<https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad>

<https://docs.microsoft.com/en-us/azure/app-service/configure-authentication-provider-aad>

upvoted 1 times

🗨️ **ytingyeu** 3 years, 4 months ago

imo managed identities is for resource-to-resource not for client-to-resource

upvoted 3 times

You are a developer for a SaaS company that offers many web services.

All web services for the company must meet the following requirements:

- ⇒ Use API Management to access the services
- ⇒ Use OpenID Connect for authentication
- ⇒ Prevent anonymous usage

A recent security audit found that several web services can be called without any authentication.

Which API Management policy should you implement?

- A. jsonp
- B. authentication-certificate
- C. check-header
- D. validate-jwt

**Suggested Answer: D**

Add the validate-jwt policy to validate the OAuth token for every incoming request.

Incorrect Answers:

A: The jsonp policy adds JSON with padding (JSONP) support to an operation or an API to allow cross-domain calls from JavaScript browser-based clients.

JSONP is a method used in JavaScript programs to request data from a server in a different domain. JSONP bypasses the limitation enforced by most web browsers where access to web pages must be in the same domain.

JSONP - Adds JSON with padding (JSONP) support to an operation or an API to allow cross-domain calls from JavaScript browser-based clients.

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

Community vote distribution

D (100%)

🗳️ **Cornholioz** Highly Voted 4 years, 7 months ago

Correct. It's validate-jwt for Open ID, API, secure authentication.

upvoted 47 times

🗳️ **mlantonis** Highly Voted 4 years, 1 month ago

Correct Answer: D

Add the validate-jwt policy to validate the OAuth token for every incoming request.

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

upvoted 18 times

🗳️ **Vichu\_1607** Most Recent 8 months, 2 weeks ago

**Selected Answer: D**

The correct answer is D. validate-jwt.

The validate-jwt policy in Azure API Management allows you to enforce validation of JSON Web Tokens (JWT) from the client as part of the API call. This can be used to ensure that the caller has authenticated with OpenID Connect, as OpenID Connect uses JWTs for its identity tokens.

Option A, jsonp, is used for wrapping JSON responses in a JSONP callback function, and does not provide any authentication.

Option B, authentication-certificate, is used for client certificate authentication, not OpenID Connect.

Option C, check-header, is used for checking the existence of certain headers, and while it could be used to check for the existence of an Authorization header, it would not validate the JWT itself.

upvoted 2 times

- 🗳️ 👤 **applepie** 1 year, 11 months ago  
got this question today, answer D - validate-jwt - 7/30/2023, score 895/1000  
upvoted 3 times

- 🗳️ 👤 **Videira** 2 years, 4 months ago  
On my exam 2023-02-25  
upvoted 2 times

- 🗳️ 👤 **sam5678** 2 years, 8 months ago  
correct  
upvoted 1 times

- 🗳️ 👤 **serpevi** 2 years, 9 months ago  
Got this in 09/22 , went with the most voted answer, score 927.  
upvoted 3 times

- 🗳️ 👤 **Eltooth** 3 years ago  
**Selected Answer: D**  
D is correct answer.  
upvoted 3 times

- 🗳️ 👤 **JotaEleEfe** 3 years, 3 months ago  
**Selected Answer: D**  
I think validate-jwt is the correct option.  
upvoted 2 times

- 🗳️ 👤 **iamstudying** 3 years, 3 months ago  
**Selected Answer: D**  
validate-jwt, buddies  
upvoted 1 times

- 🗳️ 👤 **Freidrich** 3 years, 4 months ago  
**Selected Answer: D**  
The proposed answer is correct D: validate-jwt.  
upvoted 1 times

- 🗳️ 👤 **mabusalma** 3 years, 5 months ago  
Can someone let me understand why a check-header isn't a valid answer?

<https://docs.microsoft.com/en-us/azure/api-management/api-management-access-restriction-policies#CheckHTTPHeader>

Use the check-header policy to enforce that a request has a specified HTTP header. You can optionally check to see if the header has a specific value or check for a range of allowed values. If the check fails, the policy terminates request processing and returns the HTTP status code and error message specified by the policy.  
upvoted 1 times

- 🗳️ 👤 **coffecold** 2 years, 8 months ago  
JWT is a token issued by an identity provider, sent in the request, and that can be verified by the receiving site. A HTTP header is not issued by an identity provider.  
upvoted 4 times

- 🗳️ 👤 **mcbc** 3 years, 10 months ago  
correct  
upvoted 2 times

- 🗳️ 👤 **Peter304403** 4 years ago  
The answer might be correct, but in the explanation of the incorrect answers it sounds very confusing to read jsonp described 3-times in a similar way? Since I don't know and I couldn't find anything: Could someone please explain the other two answers, what it is or where to find something about that?  
upvoted 4 times

- 🗳️ 👤 **UnknowMan** 4 years, 1 month ago  
correct  
upvoted 1 times

  **glam** 4 years, 1 month ago

correct

upvoted 2 times

  **demius\_8** 4 years, 2 months ago

Answer is correct. Source: <https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

upvoted 4 times



## DRAG DROP -

Contoso, Ltd. provides an API to customers by using Azure API Management (APIM). The API authorizes users with a JWT token.

You must implement response caching for the APIM gateway. The caching mechanism must detect the user ID of the client that accesses data for a given location and cache the response for that user ID.

You need to add the following policies to the policies file:

- ⇒ a set-variable policy to store the detected user identity
- ⇒ a cache-lookup-value policy
- ⇒ a cache-store-value policy
- ⇒ a find-and-replace policy to update the response body with the user profile information

To which policy section should you add the policies? To answer, drag the appropriate sections to the correct policies. Each section may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

### Answer Area

Policy section	Policy	Policy section
	Set-variable	<input type="text"/>
<input type="text" value="Inbound"/>	Cache-lookup-value	<input type="text"/>
<input type="text" value="Outbound"/>	Cache-store-value	<input type="text"/>
	Find-and-replace	<input type="text"/>

### Answer Area

Policy section	Policy	Policy section
	Set-variable	<input type="text" value="Inbound"/>
<input type="text" value="Inbound"/>	Cache-lookup-value	<input type="text" value="Inbound"/>
<input type="text" value="Outbound"/>	Cache-store-value	<input type="text" value="Outbound"/>
	Find-and-replace	<input type="text" value="Outbound"/>

Suggested Answer:

Box 1: Inbound.

A set-variable policy to store the detected user identity.

Example:

```
<policies>
```

```
<inbound>
```

```
<!-- How you determine user identity is application dependent -->
```

```
<set-variable
```

```
name="enduserid"
```

```
value="@ (context.Request.Headers.GetValueOrDefault("Authorization","").Split(' ')[1].AsJwt()?.Subject)" />
```

Box 2: Inbound -

A cache-lookup-value policy -

Example:

```
<inbound>
<base />
<cache-lookup vary-by-developer="true | false" vary-by-developer-groups="true | false" downstream-caching-type="none | private | public" must-revalidate="true | false">
<vary-by-query-parameter>parameter name</vary-by-query-parameter> <!-- optional, can repeated several times -->
</cache-lookup>
</inbound>
```

Box 3: Outbound -  
A cache-store-value policy.

Example:

```
<outbound>
<base />
<cache-store duration="3600" />
</outbound>
```

Box 4: Outbound -  
A find-and-replace policy to update the response body with the user profile information.

Example:

```
<outbound>
<!-- Update response body with user profile-->
<find-and-replace
from="$userprofile$"
to="@((string)context.Variables["userprofile"])" />
<base />
</outbound>
```

Reference:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-caching-policies> <https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

  **Ave**  4 years, 7 months ago

Inbound

Inbound

Inbound

Outbound

upvoted 185 times


  **1CY1** 11 months, 3 weeks ago

Inbound, Inbound, Outbound, Outbound.

Cache store S/B outbound.

<https://learn.microsoft.com/en-us/azure/api-management/cache-store-policy#examples>



upvoted 1 times

  **NetoZee** 7 months, 3 weeks ago

it's wrong cache-store is diferente cache-store-value. you can see it here

<https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

upvoted 2 times

  **russellg** 4 years, 7 months ago

how would you cache the value on an inbound request? yo don't have the value yet!

upvoted 12 times

  **Tommy202024** 4 years, 6 months ago

Look at this example

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#architecture>

upvoted 9 times

  **[Removed]** 4 years, 5 months ago

The example you provided is for caching user information from the incoming request. The question is about storing the outgoing response.



So Inbound, Inbound, Outbound, Outbound is correct.

upvoted 27 times

  **pac1311** 4 years, 4 months ago



Think it's great that he disproved his own claim :).

upvoted 6 times

  **Azprep** 3 years, 6 months ago

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

upvoted 4 times

  **Azprep** 3 years, 2 months ago

Answers are

Inbound

Inbound

Inbound



Outbound

upvoted 3 times

  **SachinV** 1 year, 7 months ago

looks correct, in Nov182023, scored 962

upvoted 6 times

  **Ciupaz** 1 year, 6 months ago

Best score ever, compliments.

upvoted 4 times

  **taupokk**  4 years, 7 months ago

The answer provided is correct.

upvoted 38 times

  **robertob** 4 years, 7 months ago

yes is correct: <policies>

<inbound>

<base />

<cache-lookup vary-by-developer="false" vary-by-developer-groups="false" downstream-caching-type="none" must-revalidate="true" caching-type="internal" >

<vary-by-query-parameter>version</vary-by-query-parameter>

</cache-lookup>

</inbound>

<outbound>



<cache-store duration="seconds" />

<base />

</outbound>

</policies>

upvoted 8 times

  **cbn** 4 years, 4 months ago

This is not cache-store-value. However the answer seems correct as per @profesorklaus answer below.

upvoted 4 times

  **cmmr**  7 months, 1 week ago



the discussion on whether the 3rd option is inbound or outbound is valid since the cache-store-value can be used in both sections. In this case we don't want to store the result but rather which user the request is from without JWT parsing again. This saves computational effort

upvoted 1 times

  **Sinae** 8 months, 1 week ago

Answer should be Inbound, Inbound, Inbound, Outbound.

upvoted 1 times

  **FeriAZ** 1 year, 5 months ago

Set-Variable Policy

Purpose: To store the detected user identity, which is likely extracted from the JWT token.

Placement: Inbound. You need to extract and store the user ID as the request comes in, before any caching logic is applied.

Cache-Lookup-Value Policy

Purpose: To check if there's a cached response available for the specific user ID.

Placement: Inbound. This should happen early in the request processing to determine if a cached response can be used instead of forwarding the request to the backend.

Cache-Store-Value Policy

Purpose: To store the response in the cache associated with the specific user ID.

Placement: Outbound. After the backend service generates a response, this policy stores it in the cache before sending it to the client.

Find-and-Replace Policy

Purpose: To update the response body with the user profile information.

Placement: Outbound. This modification should be done after the backend service has generated the response and just before it's sent back to the client.

upvoted 5 times


  **manopeydakon** 1 year, 5 months ago

The cache-store-value policy in Azure API Management (APIM) is used in the inbound policy section. This policy is typically applied before the request reaches the backend service, allowing you to cache the response based on the specified conditions or keys.

In the context of caching, the inbound policy section refers to policies that are executed before the request is sent to the backend, and the outbound policy section refers to policies that are executed before the response is sent back to the client.

So, when you use cache-store-value to cache a response, you want to do this on the inbound side, ensuring that the caching occurs before the request is forwarded to the backend service.

upvoted 2 times

  **11\_NickName\_11** 1 year, 7 months ago

Inbound

Inbound

Inbound

Outbound

<https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#fragment-caching>

upvoted 1 times

  **dddddd111** 1 year, 7 months ago

Inbound

Inbound

Outbound

Outbound

Reference: <https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#architecture>

upvoted 1 times

  **dddddd111** 1 year, 7 months ago


Sorry, the 3rd one is Inbound.

upvoted 2 times

  **JH81** 2 years ago

Got this on 6/28/2023 and passed with 850. Went with highly voted answer.



upvoted 4 times

  **NombreFalso** 2 years, 4 months ago

>> The caching mechanism must detect the user ID of the client that accesses data for a given location and \*cache the response for that user ID.\*

Thus cache store should be outbound guys

upvoted 3 times

  **Esward** 2 years, 5 months ago

Given answers are correct.

Inbound

Inbound

outbound->in general the response come in outbound section

outbound

<https://learn.microsoft.com/en-us/azure/api-management/api-management-howto-cache>

upvoted 4 times

  **kkop** 2 years, 6 months ago

Here is Answer:<https://learn.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#architecture>

inbound,  
inbound,  
inbound  
outbound  
upvoted 5 times

🗳️ 👤 **elequiel** 2 years, 8 months ago

Got it in exam 20/10/2022  
upvoted 3 times

🗳️ 👤 **Knightie** 2 years, 10 months ago

see the xml tag of inbound and outbound  
<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>  
upvoted 1 times

🗳️ 👤 **Lucky\_me** 3 years ago

```
<policies>
<inbound>
<base />
<cache-lookup vary-by-developer="false" vary-by-developer-groups="false" downstream-caching-type="none" must-revalidate="true" caching-
type="internal" >
<vary-by-query-parameter>version</vary-by-query-parameter>
</cache-lookup>
</inbound>
<outbound>
<cache-store duration="seconds" />
<base />
</outbound>
</policies>
```

upvoted 1 times

🗳️ 👤 **vavra** 3 years ago

I think it's B, because of this:  
<https://github.com/AzureAD/microsoft-authentication-library-for-dotnet/wiki/Integrated-Windows-Authentication>  
Federated users only, i.e. those created in an Active Directory and backed by Azure Active Directory. Users created directly in AAD, without AD backing  
- managed users - cannot use this auth flow. This limitation does not affect the Username/Password flow.  
upvoted 1 times

🗳️ 👤 **kjfdzkkbsm** 3 years ago

Your'e at the wrong question, mate  
upvoted 8 times

🗳️ 👤 **Azprep** 3 years, 2 months ago

Inbound  
Inbound  
Inbound  
Outbound  
upvoted 3 times

DRAG DROP -

You are developing an Azure solution.

You need to develop code to access a secret stored in Azure Key Vault.

How should you complete the code segment? To answer, drag the appropriate code segments to the correct location. Each code segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

#### Code segments

DefaultAzureCredential

ClientSecretCredential

CloudClients

SecretClient

#### Answer Area

```
string var1 = Environment.GetEnvironmentVariable("KEY_VAULT_URI");
var var2 = new Code segment ( new Uri(var1), new Code segment ());
```

#### Suggested Answer:

##### Code segments

ClientSecretCredential

CloudClients

##### Answer Area

```
string var1 = Environment.GetEnvironmentVariable("KEY_VAULT_URI");
var var2 = new SecretClient ( new Uri(var1), new DefaultAzureCredential ());
```

Box 1: SecretClient -

Box 2: DefaultAzureCredential -

In below example, the name of your key vault is expanded to the key vault URI, in the format "https://<your-key-vault-name>.vault.azure.net".

This example is using 'DefaultAzureCredential()' class from Azure Identity Library, which allows to use the same code across different environments with different options to provide identity. string keyVaultName = Environment.GetEnvironmentVariable("KEY\_VAULT\_NAME"); var kvUri = "https://" + keyVaultName + ".vault.azure.net"; var client = new SecretClient(new Uri(kvUri), new DefaultAzureCredential());

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-net>

 **mlantonis** Highly Voted 3 years ago

Answer is correct.

```
string keyVaultName = Environment.GetEnvironmentVariable("KEY_VAULT_NAME");
var kvUri = "https://" + keyVaultName + ".vault.azure.net";
var client = new SecretClient(new Uri(kvUri), new DefaultAzureCredential());
```


Box 1: SecretClient

Box 2: DefaultAzureCredential

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/secrets/quick-create-net>

upvoted 79 times

 **Esward** 1 year, 5 months ago

Agreed

upvoted 1 times

 **SachinV** 7 months, 1 week ago

Agreed looks correct, in Nov182023, scored 962  
upvoted 2 times

🗲️ 👤 **oescm** Highly Voted 2 years, 4 months ago  
Got this one 02/2022. Went with most voted  
upvoted 9 times

🗲️ 👤 **JH81** Most Recent 1 year ago  
Got this on 6/28/2023 and passed with 850. Went with answer.  
upvoted 2 times

🗲️ 👤 **Saluk\_DE** 1 year, 3 months ago  
Question was on exam 2023-03-30  
upvoted 1 times

🗲️ 👤 **mabdo** 1 year, 4 months ago  
on exam 02/23  
upvoted 1 times

🗲️ 👤 **kociuba** 1 year, 6 months ago  
got it on my exam 30-12-2022  
score: 818  
upvoted 3 times

🗲️ 👤 **adilkhan** 1 year, 5 months ago  
are all these questions enough to pass the exam?  
upvoted 2 times

🗲️ 👤 **OPT\_001122** 1 year, 6 months ago  
1: SecretClient  
2: DefaultAzureCredential  
upvoted 2 times

🗲️ 👤 **vcfvct** 1 year, 8 months ago  
Had this question on 10/21/22  
upvoted 2 times

🗲️ 👤 **Eltooth** 2 years ago  
Secret  
Default  
upvoted 2 times

🗲️ 👤 **edengoforit** 2 years, 6 months ago  
DefaultAzureCredential  
DefaultAzureCredential is appropriate for most scenarios where the application is intended to run in the Azure Cloud. This is because the DefaultAzureCredential determines the appropriate credential type based of the environment it is executing in. It supports authenticating both as a service principal or managed identity, and can be configured so that it will work both in a local development environment or when deployed to the cloud.

The DefaultAzureCredential will first attempt to authenticate using credentials provided in the environment. In a development environment you can authenticate as a service principal with the DefaultAzureCredential by providing configuration in environment variables as described in the next section.

If the environment configuration is not present or incomplete, the DefaultAzureCredential will then determine if a managed identity is available in the current environment. Authenticating as a managed identity requires no configuration, but does require platform support. See the managed identity documentation for more details on this.  
upvoted 8 times

🗲️ 👤 **ADJ85** 2 years, 10 months ago  
Correct  
upvoted 4 times

You are developing an Azure App Service REST API.

The API must be called by an Azure App Service web app. The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).

You need to configure the API to make the updates.

Which two tools should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Microsoft Graph API
- B. Microsoft Authentication Library (MSAL)
- C. Azure API Management
- D. Microsoft Azure Security Center
- E. Microsoft Azure Key Vault SDK

**Suggested Answer: AC**

A: You can use the Azure AD REST APIs in Microsoft Graph to create unique workflows between Azure AD resources and third-party services. Enterprise developers use Microsoft Graph to integrate Azure AD identity management and other services to automate administrative workflows, such as employee onboarding (and termination), profile maintenance, license deployment, and more.

C: API Management (APIM) is a way to create consistent and modern API gateways for existing back-end services.

API Management helps organizations publish APIs to external, partner, and internal developers to unlock the potential of their data and services.

Reference:

<https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview>

Community vote distribution

AB (100%)

  **Cholo981**  4 years ago

The Answer is Graph API and MSAL.

API management as nothing to do with the developing of the REST API...

upvoted 108 times

  **AndrewJB** 3 years, 7 months ago



Correct. To access Graph API user must be logged in, so MSAL is required. APIM is optional and not mandatory for this scenario

upvoted 7 times

  **Renirechner** 3 years, 3 months ago

Furthermore the API is not intended for public use but internal use by another web app. So there is no need for APIM.

upvoted 6 times

  **pazetox** 2 years, 11 months ago

MSAL can be used to provide secure access to Microsoft Graph.

<https://docs.microsoft.com/en-us/azure/active-directory/develop/msal-overview>

upvoted 2 times

  **Esward** 2 years, 5 months ago

Agreed.

<https://learn.microsoft.com/en-us/graph/auth-cloudsolutionprovider?tabs=azuread>

upvoted 2 times

  **mlantonis**  4 years ago

Answer is correct.

A: Microsoft Graph to retrieve and update user profile info from Azure AD.

C: API Management (APIM)



Reference:

<https://docs.microsoft.com/en-us/graph/azuread-identity-access-management-concept-overview>

<https://docs.microsoft.com/en-us/graph/overview>

<https://docs.microsoft.com/en-us/azure/api-management/api-management-key-concepts>

upvoted 22 times

🗨️ **ensa** 3 years, 8 months ago

Fiest time see you are wrong.

AAAAAAAABBBBBBBBBB

upvoted 10 times

🗨️ **cooyoote** 3 years, 10 months ago

So wrong, API Management is supposed to create API gateways and rules to manage API calls

upvoted 9 times

🗨️ **J\_obando13** Most Recent 1 month, 2 weeks ago

Selected Answer: AB

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 2 times

🗨️ **vizay** 1 month, 3 weeks ago

Selected Answer: AB

Microsoft Authentication Library (MSAL) is a programming library provided by Microsoft that helps applications securely sign in users and get tokens to access Microsoft APIs (like Microsoft Graph) and other Azure-protected resources.

upvoted 1 times

🗨️ **f09229d** 3 months ago

Selected Answer: AB

In exam 30/3/2025 with this answer, I scored 907/1000

upvoted 1 times

🗨️ **Vichu\_1607** 8 months ago

Selected Answer: AB

Microsoft Graph API: Used to interact with Azure AD to retrieve and update user profile information.

Microsoft Authentication Library (MSAL): Used to handle authentication and acquire tokens for accessing Microsoft Graph API.

upvoted 2 times

🗨️ **Christian\_garcia\_martin** 11 months, 1 week ago

A and B are correct.

upvoted 1 times

🗨️ **cb98160** 1 year, 3 months ago

Selected Answer: AB

You don't need a APIM.

upvoted 1 times

🗨️ **manopeydakon** 1 year, 5 months ago

To configure the Azure App Service REST API to retrieve and update user profile information stored in Azure Active Directory (Azure AD), you should use the following tools:

A. Microsoft Graph API: The Microsoft Graph API allows you to interact with data in Azure AD, including retrieving and updating user profile information.

B. Microsoft Authentication Library (MSAL): MSAL is used for handling authentication in your application. It helps you authenticate users and acquire access tokens, which are necessary when making requests to the Microsoft Graph API.

Therefore, the correct answers are A (Microsoft Graph API) and B (Microsoft Authentication Library).

upvoted 1 times

🗄️ 👤 **Neckbread** 1 year, 5 months ago

**Selected Answer: AB**

Apim??????

upvoted 1 times

🗄️ 👤 **Dixavado** 1 year, 9 months ago

**Selected Answer: AB**

It was on my exam today (2023-09-26) - score 850

upvoted 4 times

🗄️ 👤 **nikipediaa** 1 year, 9 months ago

Question was on exam 2023-09-26. The order of answers was changed.

upvoted 3 times

🗄️ 👤 **deathRac3** 2 years, 1 month ago

**Selected Answer: AB**

trythis

upvoted 1 times

🗄️ 👤 **[Removed]** 2 years, 2 months ago

Got this in exam today (5 April 2023)

upvoted 4 times

🗄️ 👤 **NombreFalso** 2 years, 4 months ago

**Selected Answer: AB**

Choosing A implies using B too. APIM is totally unrelated.

upvoted 3 times

🗄️ 👤 **Jeff8888** 2 years, 5 months ago

Graph Api, Azure APIM

upvoted 1 times

🗄️ 👤 **Michael2023** 2 years, 5 months ago

"The API must retrieve and update user profile information stored in Azure Active Directory (Azure AD).

You need to configure the API to make the updates."

this is a hint that you want programmatically update the user data stored in Azure AD...

so Graph API and MSAL is the right answer :)

upvoted 1 times

You develop a REST API. You implement a user delegation SAS token to communicate with Azure Blob storage.

The token is compromised.

You need to revoke the token.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Revoke the delegation key.
- B. Delete the stored access policy.
- C. Regenerate the account key.
- D. Remove the role assignment for the security principle.

**Suggested Answer: AB**

A: Revoke a user delegation SAS -

To revoke a user delegation SAS from the Azure CLI, call the `az storage account revoke-delegation-keys` command. This command revokes all of the user delegation keys associated with the specified storage account. Any shared access signatures associated with those keys are invalidated.


B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier. Changing the signed identifier breaks the associations between any existing signatures and the stored access policy. Deleting or renaming the stored access policy immediately effects all of the shared access signatures associated with it.

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegation-sas-create-cli.md>

<https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying-or-revoking-a-stored-access-policy>

Community vote distribution



**jay158** Highly Voted 4 years ago

Given answer is incorrect.

Correct answers are A , D

<https://docs.microsoft.com/en-us/rest/api/storageservices/create-user-delegation-sas#revoke-a-user-delegation-sas>

upvoted 113 times

**jkes80** 3 years, 11 months ago

Your link literally says you're correct it should be A and D.

"If you believe that a SAS has been compromised, then you should revoke the SAS. You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal used to create the SAS."

upvoted 22 times

**Esward** 2 years, 5 months ago

I greed with A, D

upvoted 2 times

**maukaba** 3 years, 6 months ago

There're two ways to create a SAS:

(1). The "standard" way to generate a SAS token is to use the storage account key.

(2). by using "managed identities" with a technique is called a "user delegation" SAS, and it allows you to sign the signature with Azure AD credentials instead of with the storage account key.

This question is (2) hence A, D is correct

REF: <https://markheath.net/post/user-delegation-sas>

upvoted 3 times

**ML333** 3 years ago

Answer definitely A, D as per provided link - All exam questions answers should be as straightforward as this Microsoft! So much ambiguity in 90% of the questions.

upvoted 3 times

**mlantonis** Highly Voted 4 years ago

Answer seems correct.

A: Use az storage account revoke-delegation-keys command.

B: To revoke a stored access policy, you can either delete it, or rename it by changing the signed identifier.

Reference:

<https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/storage/blobs/storage-blob-user-delegation-sas-create-cli.md>



<https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy#modifying-or-revoking-a-stored-access-policy>

[https://docs.microsoft.com/en-us/cli/azure/storage/account?view=azure-cli-latest#az\\_storage\\_account\\_revoke\\_delegation\\_keys](https://docs.microsoft.com/en-us/cli/azure/storage/account?view=azure-cli-latest#az_storage_account_revoke_delegation_keys)  
upvoted 16 times

  **Christian\_garcia\_martin** 11 months, 1 week ago

No . Note that user delegation SAS token used in the scenario does not support Stored Access Policy so B is wrong and D is NOT valid too : The role assignments are not directly linked to SAS (Shared Access Signature) tokens in Blob Storage. Revoking a role wouldn't invalidate a SAS token. So correct Answers are A and C.

upvoted 1 times

  **Azprep** 3 years, 2 months ago

Answers should be A& D

upvoted 1 times

  **Jurgen1234** 3 years, 12 months ago

Incorrect it should be D not B, see <https://docs.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy>

Stored access policies are not supported for the user delegation SAS or the account SAS..

upvoted 3 times

  **Vichu\_1607**  8 months ago

**Selected Answer: AD**

If you believe that a SAS has been compromised, you should revoke it. You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal that's used to create the SAS.

upvoted 1 times

  **4bd3116** 10 months, 1 week ago

**Selected Answer: AC**

Revoke, Regenerate

upvoted 1 times

  **prabhjot** 10 months, 3 weeks ago

May be A and C

upvoted 1 times

  **Christian\_garcia\_martin** 11 months, 1 week ago

Guis A and C . The role assignments are not directly linked to SAS (Shared Access Signature) tokens in Blob Storage. Removing a role wouldn't invalidate a SAS token

upvoted 1 times

  **manopeydakon** 1 year, 5 months ago

A,C

When dealing with a compromised user delegation SAS token in Azure Blob storage, you can take the following actions:

A. Revoke the delegation key: User delegation SAS tokens are generated based on a delegation key. Revoking the delegation key would invalidate any tokens generated with that key.

C. Regenerate the account key: Regenerating the account key would effectively invalidate all SAS tokens, including user delegation SAS tokens, associated with the storage account. This is a more drastic measure and should be carefully considered, as it affects all tokens, not just the compromised one.

Therefore, the correct answers are A (Revoke the delegation key) and C (Regenerate the account key).

upvoted 4 times

🗨️ 👤 **FeriAZ** 1 year, 5 months ago

A. Revoke the Delegation Key

Approach: By revoking the user delegation key used to create the SAS token, you can effectively invalidate the SAS token.

Feasibility: Azure Blob Storage allows you to revoke user delegation keys, which will invalidate any SAS tokens created with them.

Effectiveness: This is a direct way to revoke a compromised User Delegation SAS token.

D. Remove the Role Assignment for the Security Principle

Approach: This involves removing the Azure AD role assignment that grants permissions to the user or service principal associated with the SAS token.

Feasibility: By removing or altering the role assignment in Azure AD, you can effectively revoke access permissions that the SAS token grants.

Effectiveness: This can be an effective way to revoke access, though it may be broader than just invalidating a specific SAS token.

upvoted 1 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

A and D!

upvoted 1 times

🗨️ 👤 **nikipediaa** 1 year, 9 months ago

Question was on exam 2023-09-26

upvoted 1 times

🗨️ 👤 **Teit\_examtopics** 1 year, 9 months ago

If you believe that a SAS has been compromised, you should revoke it. You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal that's used to create the SAS.

<https://learn.microsoft.com/en-us/rest/api/storageservices/create-user-delegation-sas>

upvoted 1 times

🗨️ 👤 **Ivanlco** 1 year, 10 months ago

C By regenerating the you make the prevues key unusable why isn't that right

upvoted 1 times

🗨️ 👤 **RKyawal** 1 year, 12 months ago

<https://learn.microsoft.com/en-us/rest/api/storageservices/define-stored-access-policy>

cannot be B as per this link..not supported

upvoted 1 times

🗨️ 👤 **DP\_Bhatt** 2 years, 5 months ago

**Selected Answer: AD**

Correct answers A and D

upvoted 1 times

🗨️ 👤 **rosa\_0617** 2 years, 10 months ago

**Selected Answer: AD**

If you believe that a SAS has been compromised, then you should revoke the SAS.

You can revoke a user delegation SAS either by revoking the user delegation key, or by changing or removing RBAC role assignments for the security principal used to create the SAS.

upvoted 4 times

🗨️ 👤 **nuiowmokka** 2 years, 10 months ago

Can we pass the exam just by studying questions tilll Page 22? I don't have contributor access and can't afford right now.

upvoted 5 times

🗨️ 👤 **ASAAAAS2** 2 years, 11 months ago

**Selected Answer: AD**

A & D for sure

upvoted 2 times

## DRAG DROP -

You are developing an Azure-hosted application that must use an on-premises hardware security module (HSM) key.

The key must be transferred to your existing Azure Key Vault by using the Bring Your Own Key (BYOK) process.

You need to securely transfer the key to Azure Key Vault.

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.

Select and Place:

**Actions**

Generate a key transfer blob file by using the HSM vendor-provided tool.

Generate a Key Exchange Key (KEK).

Create a custom policy definition in Azure Policy.

Run the `az keyvault key import` command.

Run the `az keyvault key restore` command.

Retrieve the Key Exchange Key (KEK) public key.

**Answer Area****Suggested Answer:****Actions**

Create a custom policy definition in Azure Policy.

Run the `az keyvault key restore` command.

**Answer Area**

Generate a Key Exchange Key (KEK).

Retrieve the Key Exchange Key (KEK) public key.

Generate a key transfer blob file by using the HSM vendor-provided tool.

Run the `az keyvault key import` command.

To perform a key transfer, a user performs following steps:

- ⇒ Generate KEK.
- ⇒ Retrieve the public key of the KEK.
- ⇒ Using HSM vendor provided BYOK tool - Import the KEK into the target HSM and exports the Target Key protected by the KEK.
- ⇒ Import the protected Target Key to Azure Key Vault.

Step 1: Generate a Key Exchange Key (KEK).

Step 2: Retrieve the Key Exchange Key (KEK) public key.

Step 3: Generate a key transfer blob file by using the HSM vendor-provided tool.

Generate key transfer blob using HSM vendor provided BYOK tool

Step 4: Run the `az keyvault key import` command

Upload key transfer blob to import HSM-key.

Customer will transfer the Key Transfer Blob (".byok" file) to an online workstation and then run a `az keyvault key import` command to import this blob as a new

HSM-backed key into Key Vault.

To import an RSA key use this command:

`az keyvault key import`

Reference:

<https://docs.microsoft.com/en-us/azure/key-vault/keys/byok-specification>

 **aradice**  2 years, 12 months ago

correct: <https://docs.microsoft.com/en-us/azure/key-vault/keys/byok-specification#user-steps>

upvoted 52 times

 **Esward** 1 year, 5 months ago

Agreed

upvoted 1 times

🗨️ 👤 **Esward** 1 year, 5 months ago

Given answers are correct  
upvoted 1 times

🗨️ 👤 **TakumaK** 2 years, 12 months ago

agree with you!  
upvoted 4 times

🗨️ 👤 **MasterQuestMaster** Highly Voted 🏆 2 years, 5 months ago

Got this on the exam. :)  
top kek  
upvoted 23 times

🗨️ 👤 **mandynotmandy** 2 years, 3 months ago

this is still showing up in feb 2022 exams, the kek definitely made this question unforgettable for me  
top kek  
upvoted 8 times

🗨️ 👤 **vizay** Most Recent 🔥 1 month, 3 weeks ago

You encrypt your key using Azure's public key ➡️ make a secure file ➡️ upload it to Azure ➡️ Azure decrypts and stores it.  
upvoted 1 times

🗨️ 👤 **kotireddy4120** 7 months, 1 week ago

Given Answer is correct  
upvoted 1 times

🗨️ 👤 **mewan** 1 year, 1 month ago

is it enough to pass the exam to read up to 22 pages  
upvoted 7 times

🗨️ 👤 **Videira** 1 year, 4 months ago

On my exam 2023-02-25  
upvoted 2 times

🗨️ 👤 **NombreFalso** 1 year, 4 months ago

based kek question  
upvoted 2 times

🗨️ 👤 **Jeff8888** 1 year, 5 months ago

Agreed  
upvoted 1 times

🗨️ 👤 **petitbilly** 2 years, 3 months ago

Got it in exam 03/22  
upvoted 4 times

🗨️ 👤 **oescm** 2 years, 4 months ago

Got this one 02/2022. Went with most voted  
upvoted 6 times

🗨️ 👤 **lugospod** 2 years, 5 months ago

Got this one 01/2022. Went with originally proposed solution  
upvoted 6 times

🗨️ 👤 **danila16030** 2 years, 5 months ago

Ho many questions from you test you find on this site? Just interesting)  
upvoted 1 times

🗨️ 👤 **mcbc** 2 years, 10 months ago

Generate KEK.  
Retrieve the public key of the KEK.  
Using HSM vendor provided BYOK tool - Import the KEK into the target HSM and exports the Target Key protected by the KEK.  
Import the protected Target Key to Azure Key Vault.  
upvoted 2 times

You develop and deploy an Azure Logic app that calls an Azure Function app. The Azure Function app includes an OpenAPI (Swagger) definition and uses an Azure Blob storage account. All resources are secured by using Azure Active Directory (Azure AD). The Azure Logic app must securely access the Azure Blob storage account. Azure AD resources must remain if the Azure Logic app is deleted. You need to secure the Azure Logic app. What should you do?

- A. Create a user-assigned managed identity and assign role-based access controls.
- B. Create an Azure AD custom role and assign the role to the Azure Blob storage account.
- C. Create an Azure Key Vault and issue a client certificate.
- D. Create a system-assigned managed identity and issue a client certificate.
- E. Create an Azure AD custom role and assign role-based access controls.

#### Suggested Answer: A

To give a managed identity access to an Azure resource, you need to add a role to the target resource for that identity.

Note: To easily authenticate access to other resources that are protected by Azure Active Directory (Azure AD) without having to sign in and provide credentials or secrets, your logic app can use a managed identity (formerly known as Managed Service Identity or MSI). Azure manages this identity for you and helps secure your credentials because you don't have to provide or rotate secrets.

If you set up your logic app to use the system-assigned identity or a manually created, user-assigned identity, the function in your logic app can also use that same identity for authentication.

Reference:


<https://docs.microsoft.com/en-us/azure/logic-apps/create-managed-service-identity> <https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-mutual-certificates-for-clients>

Community vote distribution


A (100%)

 **aradice** Highly Voted 3 years, 12 months ago

correct "Azure AD resources must remain if the Azure Logic app is deleted."  
upvoted 47 times

 **TonyMel** 2 years, 3 months ago

correct, in 2023Mar24, score: 904/1000.  
upvoted 7 times

 **debanjan10** Highly Voted 3 years, 8 months ago


User-assigned keys are individual components. Even if the logic apps are deleted, the keys remain. But in case of system-assigned keys, those are auto generated and are deleted when the Azure resources themselves are deleted.  
upvoted 23 times

 **Vichu\_1607** Most Recent 8 months ago

Selected Answer: A

A. Create a user-assigned managed identity and assign role-based access controls.

To securely access the Azure Blob storage account from the Azure Logic app while ensuring that Azure AD resources remain if the Azure Logic app is deleted, you should use a user-assigned managed identity. This approach allows you to manage the identity independently of the Logic app's lifecycle.  
upvoted 2 times

 **Ciupaz** 1 year, 5 months ago

Selected Answer: A

Azure AD resources must remain if the Azure Logic app is deleted -> User Managed Identity  
upvoted 1 times

 **FeriAZ** 1 year, 5 months ago



A. Create a user-assigned managed identity and assign role-based access controls.

This approach ensures that the Azure Logic App can securely access the Azure Blob Storage with the appropriate permissions, and the managed identity remains operational and intact, independent of the Logic App's lifecycle.

upvoted 1 times

🗲️ 👤 **bgbgvf** 1 year, 6 months ago

correct answer

upvoted 1 times

🗲️ 👤 **applepie** 1 year, 11 months ago

got this question today, answer A - 7/30/2023, score 895/1000

upvoted 5 times

🗲️ 👤 **databasejamdown** 2 years, 1 month ago

**Selected Answer: A**

User assigned will persist after resource is removed

upvoted 5 times

🗲️ 👤 **RaghavMGupta** 2 years, 2 months ago

Why are other options wrong?

upvoted 1 times

🗲️ 👤 **Sulzirsha** 2 years, 4 months ago

Was on exam feb21

upvoted 1 times

🗲️ 👤 **rotimislav** 2 years, 6 months ago

Why not E. Create Azure AD custom role?

upvoted 1 times

🗲️ 👤 **Christian\_garcia\_martin** 11 months, 1 week ago

Because a custom role only cover Authorization part , option A cover authorization and authentication to protect the login app

upvoted 1 times

🗲️ 👤 **r3verse** 2 years, 7 months ago

"You need to secure the Azure Logic app.", lol, but we aren't securing the logic app, we are providing a mechanism for the logic app to securely access other resources. (a key not a lock). The logic app itself can still be accessed anonymously.

upvoted 4 times

🗲️ 👤 **EmnaDa** 2 years, 7 months ago

correct "Azure AD resources must remain if the Azure Logic app is deleted."

upvoted 1 times

🗲️ 👤 **sam5678** 2 years, 8 months ago

correct

upvoted 1 times

🗲️ 👤 **iamstudying** 3 years, 3 months ago

**Selected Answer: A**

A, buddies. Agree with @dejanban10

upvoted 2 times

🗲️ 👤 **meoukg** 3 years, 3 months ago

Got it on 03/2022, I chose A. Create a user-assigned managed identity and assign role-based access controls.

upvoted 3 times

🗲️ 👤 **mattvasc** 3 years, 4 months ago

**Selected Answer: A**

In here <https://docs.microsoft.com/en-us/azure/active-directory/managed-identities-azure-resources/overview>

We can find:

"Common use cases: Workloads where resources are recycled frequently, but permissions should stay consistent."

upvoted 1 times

## HOTSPOT -

You are developing an application that uses a premium block blob storage account. You are optimizing costs by automating Azure Blob Storage access tiers.

You apply the following policy rules to the storage account. You must determine the implications of applying the rules to the data. (Line numbers are included for reference only.)

```

01 {
02   "rules":
03   {
04     "name": "agingDataRule",
05     "enabled": true,
06     "type": "Lifecycle",
07     "definition": {
08       "filters": {
09         "blobTypes": [ "blockBlob" ],
10         "prefixMatch": [ "container1/salesorders", "container2/inventory" ]
11       },
12       "actions": {
13         "baseBlob": {
14           "tierToCool": { "daysAfterModificationGreaterThan": 60 },
15           "tierToArchive": { "daysAfterModificationGreaterThan": 120 }
16         }
17       }
18     },
19   },
20   {
21     "enabled": true,
22     "name": "lastAccessedDataRule",
23     "type": "Lifecycle",
24     "definition": {
25       "actions": {
26         "baseBlob": {
27           "enableAutoTierToHotFromCool": true,
28           "tierToCool": {
29             "daysAfterLastAccessTimeGreaterThan": 30
30           }
31         }
32       },
33       "filters": {
34         "blobTypes": [ "blockBlob" ]
35       }
36     },
37   },
38   {
39     "rules": [
40       {
41         "name": "expirationDataRule",
42         "enabled": true,
43         "type": "Lifecycle",
44         "definition": {
45           "filters": {
46             "blobTypes": [ "blockBlob" ]
47           },
48           "actions": {
49             "baseBlob": {
50               "delete": { "daysAfterModificationGreaterThan": 730 }
51             }
52           }
53         }
54       }
55     ]
56   }
57 }
58 }

```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

	Yes	No
Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.	<input type="radio"/>	<input type="radio"/>
Blobs are moved to cool storage if they have not been accessed for 30 days.	<input type="radio"/>	<input type="radio"/>
Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.	<input type="radio"/>	<input type="radio"/>
All block blobs older than 730 days will be deleted.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

Yes No

Block blobs prefixed with container1/salesorders or container2/inventory which have not been modified in over 60 days are moved to cool storage. Blobs that have not been modified in 120 days are moved to the archive tier.

☒☐

Blobs are moved to cool storage if they have not been accessed for 30 days.

☒☐

Blobs will automatically be tiered from cool back to hot if accessed again after being tiered to cool.

☒☐

All block blobs older than 730 days will be deleted.

☒☐

Box 1: Yes -

```
"rules":
{
  "name": "agingDataRule",
  "enabled": true,
  "type": "Lifecycle",
  "definition": {
    "filters": {
      "blobTypes": [ "blockBlob" ],
      "prefixMatch": [ "container1/salesorders", "container2/inventory" ]
    },
    "actions": {
      "baseBlob": {
        "tierToCool": { "daysAfterModificationGreaterThan": 60 },
        "tierToArchive": { "daysAfterModificationGreaterThan": 120 }
      }
    }
  }
}
```

Box 2: Yes -

```
"enabled": true,
"name": "lastAccessedDataRule",
"type": "Lifecycle",
"definition": {
  "actions": {
    "baseBlob": {
      "enableAutoTierToHotFromCool": true,
      "tierToCool": {
        "daysAfterLastAccessTimeGreaterThan": 30
      }
    }
  }
}
```

Box 3: Yes -

Box 4: Yes -

```
"rules": [
{
  "name": "expirationDataRule",
  "enabled": true,
  "type": "Lifecycle",
  "definition": {
    "filters": {
      "blobTypes": [ "blockBlob" ]
    },
    "actions": {
      "baseBlob": {
        "delete": { "daysAfterModificationGreaterThan": 730 }
      }
    }
  }
}
```

 **finnishr** Highly Voted 3 years, 10 months ago

Correct answer:

1. Yes
2. Yes
3. Yes
4. No

upvoted 94 times

 **vizay** 1 month, 3 weeks ago

Lifecycle management policies are supported for block blobs and append blobs in general-purpose v2, premium block blob, and Blob Storage accounts. Lifecycle management doesn't affect system containers such as the \$logs or \$web containers.

upvoted 1 times

 **[Removed]** 2 years, 4 months ago

I think you missed the difference between Blobs and Block blobs.

1. Yes
2. No - configuration is for Block Blobs and not for blob
3. No - configuration is for Block Blobs and not for blob
4. Yes - because of delete property

upvoted 1 times

  **warchoon** 2 years, 4 months ago



So 1 must be No too because of "Blobs" for 120 days

upvoted 2 times

  **TOM101** 2 years, 3 months ago



The scenario is about premium block blobs, so blobs can only be block blobs here

upvoted 3 times

  **examprepau** 8 months, 2 weeks ago

Disagree: Append Blobs can also go in Premium Block Blobs

upvoted 1 times

  **examprepau** 8 months, 2 weeks ago

So for me its

Yes: 60 Days( BlockBlobs with container1/salesorder or container2/salesorder)

NO : 120 Days (Mentions Blobs in general. where append blobs will not move as rule is for blockBlob)

NO: 30 Days (Mentions Blobs in general . where append blobs will not move as rule is for blockBlob)

NO: Auto move to hot (Mentions Blobs in general . where append blobs will not move as rule is for blockBlob)

NO: 730 (rule is daysAfterModified not all blockBlobs)



Reference:

[learn.microsoft.com/en-us/azure/storage/common/storage-account-overview](https://learn.microsoft.com/en-us/azure/storage/common/storage-account-overview)

"Premium block blobs3 Blob Storage (including Data Lake Storage1)

Premium storage account type for block blobs and append blobs. Recommended for scenarios with high transaction rates or that use smaller objects or require consistently low storage latency. Learn more about example workloads."


upvoted 1 times

  **examprepau** 8 months, 2 weeks ago

Justification/correction for above: Premium Block Blob account can have append blobs as well, all the statements except first and last says blob and not blockBlob, where all the filters are related to blockBlob. Since the first and last statement mentions blockBlob explicitly that mean it needs to be considered.

YES 4: auto move to hot as only block blobs are going to cool so that means all the blobs in cool will go to hot

upvoted 1 times

  **Esward** 2 years, 5 months ago



Agree with you!

upvoted 1 times

  **windflower555** 3 years, 9 months ago

4 No, correct answer is: should be block blobs will be deleted 730 days after last modified  
(missed deleted part in my previous comment)

upvoted 13 times

  **Vady98** 3 years, 8 months ago

Agree

I think

upvoted 1 times

  **jay158**  3 years, 12 months ago

Answer is correct



1. Yes
2. Yes
3. Yes
4. Yes

upvoted 39 times

  **vtomy** 3 years, 12 months ago

Answer is correct

upvoted 4 times


  **ucsdmiami2020** 3 years, 7 months ago

4 is Yes. Per the Microsoft docs "Some data is expected to expire days or months after creation. You can configure a lifecycle management policy to expire data by deletion based on data age. The following example shows a policy that deletes all block blobs older than 365 days."

```
"actions": {  
  "baseBlob": {  
    "delete": { "daysAfterModificationGreaterThan": 365 }  
  }  
}
```



<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier>

upvoted 3 times

  **KeithSh** 11 months, 1 week ago

It actually says "The following example shows a policy that deletes all block blobs that haven't been MODIFIED in the last 365 days." Not "older than".

upvoted 2 times

  **KeithSh** 11 months, 1 week ago

So should be NO for the 4th Q.

upvoted 2 times

  **mattvasc** 3 years, 3 months ago

I have opened an PR, and now, the docs is correct:

"The following example shows a policy that deletes all block blobs that have not been modified in the last 365 days."

PR:

<https://github.com/MicrosoftDocs/azure-docs/pull/89203>

Docs:

<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#expire-data-based-on-age>

upvoted 8 times

  **gmishra88** 2 years, 8 months ago

Bad that you cannot start a PR on this question from Microsoft :)



upvoted 1 times

  **yibuqian** 3 years, 3 months ago

#2 rule conflict with #1, so #1ans - No

#4ans is No, because the rule is not for created date

upvoted 1 times

  **wangga** 3 years, 3 months ago

4. No, It should be daysAfterCreationGreaterThan



<https://docs.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal>

upvoted 5 times

  **wafa\_chaari**  5 months ago

yes yes yes no

upvoted 1 times

  **4bd3116** 10 months, 1 week ago

Yes,Yes,No,Yes

upvoted 1 times

  **blpiek21** 1 year, 3 months ago

So confused, Premium Block Blob storage does not support tiering. Or am I missing something?

upvoted 8 times

  **cmmr** 6 months, 3 weeks ago

<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

it seems it does now

upvoted 1 times

🗨️ 👤 **SSR999** 1 year, 4 months ago

All No, as premium block blob wont support lifecycle management policies (Tiers) and access tiers

<https://learn.microsoft.com/en-us/azure/storage/blobs/storage-feature-support-in-storage-accounts#premium-block-blob-accounts>

upvoted 4 times

🗨️ 👤 **Veeresh114** 1 year, 5 months ago

The baseBlob element in a lifecycle management policy refers to the current version of a blob. The version element refers to a previous version.

upvoted 1 times

🗨️ 👤 **manopeydakon** 1 year, 5 months ago

NO,No,No,No - Tiering is not yet supported in a premium block blob storage account. For all other accounts, tiering is allowed only on block blobs and not for append and page blobs.

See: <https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

upvoted 4 times

🗨️ 👤 **CarlosTheBoldest** 1 year, 6 months ago

I got this question on my exam, 2023Dec, go with what I remember was the most voted answer. Score 902, most of the questions were here, slightly different on wording because the Azure Ad <-> Entra Id change. Case was City Power & Light. Good luck!

Important tip, you have access to microsoft learn during the exam!

upvoted 3 times

🗨️ 👤 **ww** 1 year, 7 months ago

Received this question on my test dated 28-11-2023,

went with

No,

No,

No,

No

Because as mentioned in the discussions here, Premium block blobs don't support tiering as of Nov-23.

upvoted 10 times

🗨️ 👤 **cmmr** 6 months, 3 weeks ago

they do now: Lifecycle management policies are supported for block blobs and append blobs in general-purpose v2, premium block blob, and Blob Storage accounts.

<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

upvoted 2 times

🗨️ 👤 **dddddd111** 1 year, 7 months ago

As of Oct2023, premium block blob storage account is not yet supported for tiering.

This is a tricky question, just because you see a code doesn't mean it will work. All answers are "No", unless the premium block blob storage account becomes available in the future. (If that's happen, I believe No.4 would still be no, due to the delete rule requirements)

Reference: Search the note for premium block blob storage account.

<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

upvoted 6 times

🗨️ 👤 **nanguer** 1 year, 8 months ago

All of the answers should be "No" according to this paragraph in the docs:

"Data stored in a premium block blob storage account cannot be tiered to hot, cool, cold or archive by using Set Blob Tier or using Azure Blob Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the hot tier in a different account using the Put Block From URL API or a version of AzCopy that supports this API"

<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview#blob-lifecycle-management>

upvoted 9 times

🗨️ 👤 **dddddd111** 1 year, 7 months ago

Indeed. It was also mentioned in this site <https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview>

Tiering is not yet supported in a premium block blob storage account. For all other accounts, tiering is allowed only on block blobs and not for append and page blobs.

All answers are "NO"

upvoted 4 times

🗄️ 👤 **pandugadu009** 1 year, 9 months ago

Got it today and went with

No

No

No

Yes

Passed the exam with 749 score.

upvoted 3 times

🗄️ 👤 **leviatas** 1 year, 9 months ago

Got this question in exam 2023.09.25. Contoso case Study

upvoted 1 times

🗄️ 👤 **SangeethaRamasamy** 1 year, 9 months ago

Got this in 14/09/2023 exam

upvoted 2 times

🗄️ 👤 **Elbanna** 1 year, 9 months ago

which answers did you choose?

upvoted 1 times

🗄️ 👤 **Dreamer999** 2 years ago

Answer should be : No, No ,No, No

The key is that this is premium storage account and moving between tiers are not supported using lifecycle management policy.

I tested it by creating a premium storage account with the blob type as block blob. In the life cycle management window for adding rules in the portal,

I only see option to delete the blob. No option to move to different tiers.

upvoted 11 times

🗄️ 👤 **JogSte** 1 year, 10 months ago

I tested too and Premium Block Blob Storage Account only accept Delete lifecycle, not tierToCool neither tierToArchive.

The documentation says 1-NO, 2-NO, 3-NO

"Data stored in a premium block blob storage account cannot be tiered to hot, cool, cold or archive by using Set Blob Tier or using Azure Blob

Storage lifecycle management. To move data, you must synchronously copy blobs from the block blob storage account to the hot tier in a different account using the Put Block From URL API or a version of AzCopy that supports this API."

<https://learn.microsoft.com/en-us/azure/storage/blobs/access-tiers-overview>

upvoted 2 times

🗄️ 👤 **prakash007** 2 years, 1 month ago

Correct answer :

1. Yes

2. Yes

3. Yes

4. No.

Reference: -

<https://learn.microsoft.com/en-us/azure/storage/blobs/lifecycle-management-overview?tabs=azure-portal#move-aging-data-to-a-cooler-tier>

upvoted 1 times

You are developing a solution that will use a multi-partitioned Azure Cosmos DB database. You plan to use the latest Azure Cosmos DB SDK for development.

The solution must meet the following requirements:

- ⇒ Send insert and update operations to an Azure Blob storage account.
- ⇒ Process changes to all partitions immediately.
- ⇒ Allow parallelization of change processing.

You need to process the Azure Cosmos DB operations.

What are two possible ways to achieve this goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Create an Azure App Service API and implement the change feed estimator of the SDK. Scale the API by using multiple Azure App Service instances.
- B. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK.
- C. Create an Azure Function to use a trigger for Azure Cosmos DB. Configure the trigger to connect to the container.
- D. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

#### Suggested Answer: AC

Azure Functions is the simplest option if you are just getting started using the change feed. Due to its simplicity, it is also the recommended option for most change feed use cases. When you create an Azure Functions trigger for Azure Cosmos DB, you select the container to connect, and the Azure Function gets triggered whenever there is a change in the container. Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions.

Note: You can work with change feed using the following options:

- ⇒ Using change feed with Azure Functions
- ⇒ Using change feed with change feed processor

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed>

Community vote distribution



**Zulhin** Highly Voted 3 years, 11 months ago

Answer C & D.

"What are two possible ways to achieve this goal?"

upvoted 43 times

**Chiboy** 3 years, 4 months ago

C is the correct answer. Each answer has two possible ways that in combination, gives the solution. Moreover, because of the requirement to "Process changes to all partitions immediately", D cannot be the answer or part of it since it "processes the change feed by using the pull model on the container"

upvoted 4 times

**ProtossOR89144** 11 months, 4 weeks ago

D is good - <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/read-change-feed> - it says: "However, there are some scenarios where you might want the additional low level control of the pull model. These include:

Reading changes from a particular partition key"

upvoted 1 times

**Baskman** 3 years, 4 months ago

Agree with C & D

"NOTE: Each correct selection is worth one point."

upvoted 3 times

**aradice** Highly Voted 3 years, 12 months ago



c and d ? <https://docs.microsoft.com/en-us/azure/cosmos-db/read-change-feed#azure-functions>  
<https://docs.microsoft.com/en-us/azure/cosmos-db/change-feed-pull-model#using-feedrange-for-parallelization>  
upvoted 16 times

  **Jurgen1234** 3 years, 12 months ago

I agree

upvoted 3 times

  **Vichu\_1607** Most Recent 8 months ago

Selected Answer: CD

Azure Function with Cosmos DB Trigger: Automatically handles partitioning and parallelization, making it easier to process changes in real-time.  
Azure Function with FeedIterator and FeedRange: Provides manual control over change feed processing and allows for efficient parallelization.

upvoted 1 times

  **obllew** 9 months ago

FeedIterator is meant to be implemented in a persistent process; it runs indefinitely. A serverless environment like Functions isn't really appropriate. It should be running e.g. as a daemon that auto starts and restarts on failure etc. Which makes me question if D.

upvoted 1 times

  **Occ50bf** 10 months, 3 weeks ago

It is not D. It says that it uses the pull model on the \*container\*, but a change processor should pull from Cosmos, not the storage container. Good answers are B and C.

upvoted 1 times

  **Munwalinwali** 12 months ago

Selected Answer: CD

Answer C & D.

upvoted 1 times

  **carlosfaria82** 1 year, 2 months ago

Selected Answer: CD

C: "Because Azure Functions uses the change feed processor behind the scenes, it automatically parallelizes change processing across your container's partitions."

D: "You can use the change feed pull model to consume the Azure Cosmos DB change feed at your own pace. Similar to the change feed processor, you can use the change feed pull model to parallelize the processing of changes across multiple change feed consumers."

upvoted 1 times

  **raymond\_abcd** 1 year, 5 months ago

Answer C and D

C you need to implement the lease container in the function for dynamic scaling and works with multiple partitions. See:  
<https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/change-feed-functions>.

D is also correct as it is described here: <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/read-change-feed>

upvoted 1 times

  **FeriAZ** 1 year, 5 months ago

Selected Answer: CD

C. Azure Function with Cosmos DB Trigger

Azure Functions can be triggered by changes in Cosmos DB, allowing for immediate processing of inserts and updates.

D. Azure Function with FeedIterator and FeedRange for Change Feed

This method manually handles the change feed processing, allowing for detailed control over parallelization and processing logic.

upvoted 1 times

  **manopeydakon** 1 year, 5 months ago

Both options B and D can be used to process Azure Cosmos DB operations efficiently:

B. Create a background job in an Azure Kubernetes Service and implement the change feed feature of the SDK. This option involves creating a background job using Azure Kubernetes Service (AKS) and implementing the change feed feature of the Azure Cosmos DB SDK. AKS provides scalability, and you can parallelize the processing by distributing work across multiple pods in the AKS cluster. D. Create an Azure Function that uses a FeedIterator object that processes the change feed by using the pull model on the container. Use a FeedRange object to parallelize the processing of the change feed across multiple functions.

This option involves using Azure Functions with the change feed trigger for Azure Cosmos DB. By using the FeedIterator and FeedRange objects, you can parallelize the processing of the change feed across multiple Azure Functions.

upvoted 1 times

🗨️ 👤 **Iednari** 1 year, 6 months ago

B and C for me. AKS containers can scale and job just runs in background. Funcs can scale and trigger on cosmos.

A is an api so whats it triggering on?

D also doesnt mention what it would trigger on

upvoted 1 times

🗨️ 👤 **Dianahu** 1 year, 11 months ago

isn't the change feed estimator only for monitoring? it is the change feed processor who does the work <https://learn.microsoft.com/en-us/azure/cosmos-db/nosql/how-to-use-change-feed-estimator?tabs=dotnet> so it is C&D

upvoted 2 times

🗨️ 👤 **adilkhan** 2 years, 2 months ago

C,D chatGPT

upvoted 2 times

🗨️ 👤 **Ciupaz** 1 year, 9 months ago

In my case chatGPT response is A and C.

upvoted 1 times

🗨️ 👤 **Ciupaz** 1 year, 7 months ago

Bard, instead, returns: A - D

upvoted 1 times

🗨️ 👤 **Mattt** 8 months, 1 week ago

My B,C

upvoted 1 times

🗨️ 👤 **adilkhan** 2 years, 3 months ago

CD (36%)

D (21%)

A (21%)

Other ...OMG :D

upvoted 7 times

🗨️ 👤 **giuliohome** 12 months ago

It's C and D, you can exclude A because it's sequential but in the question it's required to allow parallelization instead.

upvoted 1 times

🗨️ 👤 **mabdo** 2 years, 4 months ago

similar was on 2/23

upvoted 1 times

🗨️ 👤 **alexein74** 2 years, 5 months ago

**Selected Answer: CD**

I agree with Zulhin.

upvoted 1 times

🗨️ 👤 **Idenouter** 2 years, 5 months ago

**Selected Answer: CD**

I agree with Zulhin.

upvoted 1 times

## HOTSPOT -

You have an Azure Web app that uses Cosmos DB as a data store. You create a CosmosDB container by running the following PowerShell script:

```
$resourceGroupName = "testResourceGroup"
$accountName = "testCosmosAccount"
$databaseName = "testDatabase"
$containerName = "testContainer"
$partitionKeyPath = "/EmployeeId"
$autoscaleMaxThroughput = 5000
```

## New-AzCosmosDBSqlContainer -

```
-ResourceGroupName $resourceGroupName
-AccountName $accountName
-DatabaseName $databaseName
-Name $containerName
-PartitionKeyKind Hash
-PartitionKeyPath $partitionKeyPath
-AutoscaleMaxThroughput $autoscaleMaxThroughput
```

You create the following queries that target the container:

```
SELECT * FROM c WHERE c.EmployeeId > '12345'
SELECT * FROM c WHERE c.UserId = '12345'
```

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

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

	Yes	No
The minimum throughput for the container is 400 R/Us.	<input type="radio"/>	<input type="radio"/>
The first query statement is an in-partition query.	<input type="radio"/>	<input type="radio"/>
The second query statement is a cross-partition query.	<input type="radio"/>	<input type="radio"/>

#### Suggested Answer:

### Answer Area

	Yes	No
The minimum throughput for the container is 400 R/Us.	<input type="radio"/>	<input checked="" type="radio"/>
The first query statement is an in-partition query.	<input type="radio"/>	<input checked="" type="radio"/>
The second query statement is a cross-partition query.	<input checked="" type="radio"/>	<input type="radio"/>

Box 1: No -

You set the highest, or maximum RU/s T<sub>max</sub> you don't want the system to exceed. The system automatically scales the throughput T such that  $0.1 * T_{max} \leq T \leq$

T<sub>max</sub>.

In this example we have autoscaleMaxThroughput = 5000, so the minimum throughput for the container is 500 R/Us.

Box 2: No -

First query: SELECT \* FROM c WHERE c.EmployeeId > '12345'

Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key:

```
SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'
```

Box 3: Yes -

Example of In-partition query:

Consider the below query with an equality filter on DeviceId. If we run this query on a container partitioned on DeviceId, this query will filter to a single physical partition.

```
SELECT * FROM c WHERE c.DeviceId = 'XMS-0001'
```

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer> <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container>

 **aradice** Highly Voted 3 years, 6 months ago

correct:

1° no => line 6

2° no =>

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer#overview-of-provisioned-throughput-types>

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container#in-partition-query>

"Here's a query that has a range filter on the partition key and won't be scoped to a single physical partition. In order to be an in-partition query, the query must have an equality filter that includes the partition key: `SELECT * FROM c WHERE c.DeviceId > 'XMS-0001'`"

3° : yes => partition key is EmployeeId. <https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-query-container#cross-partition-query>


upvoted 44 times

 **xRiot007** 2 years, 9 months ago

The documentation on the second point could be better.

They put the most important part at the end and that is that even if the query uses the partition key it MUST use an equality filter to be considered in-partition.

upvoted 3 times

 **finnishr** 3 years, 4 months ago

100% correct!

upvoted 2 times

 **HimanshuNankani** Highly Voted 3 years, 2 months ago

1. No : Because as question specifies max RUs = 5000, so minimum RUs =  $5000/10 = 500$

<https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer#overview-of-provisioned-throughput-types>

2. No : In-partition query needs equality filter that is missing here

3. Yes: Though we have equality filter here, 'UserId' is not a partition key here.

<https://docs.microsoft.com/en-us/azure/cosmos-db/sql/how-to-query-container>


upvoted 27 times

 **neelkanths** Most Recent 8 months ago

Got it on 20 April 2024...Marks > 900...All questions from examtopics 400 questions...

answer is correct...

upvoted 2 times

 **FeriAZ** 11 months, 2 weeks ago

Minimum Throughput for the Container is 400 R/Us:

Answer: No. In Cosmos DB, the minimum throughput for a container with autoscale is 10% of the maximum throughput. Since the maximum is set to 5000 RUs, the minimum would be 500 RUs (10% of 5000).

The First Query Statement is an In-Partition Query:

Answer: No. Though it filters on the partition key, it uses a range condition, potentially involving multiple partitions.

The Second Query Statement is a Cross-Partition Query:

Answer: Yes. It filters on a non-partition key (UserID), necessitating a scan across multiple partitions.

upvoted 5 times

 **bgbgvfvf** 1 year ago

Given answers are correct

upvoted 1 times

 **lednari** 1 year ago

Y, N, Y

Min is 400 RUs. ( $5000/100 = 50$  which is smaller than 400)

<https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits#minimum-throughput-on-container>

upvoted 2 times

 **leviatas** 1 year, 3 months ago

Got this question in my exam 2023.09.25. Had contoso Case Study  
upvoted 2 times

🗨️ 👤 **applepie** 1 year, 5 months ago  
got this question today, answer NNY- 7/30/2023, score 895/1000  
upvoted 2 times

🗨️ 👤 **NightshadeRC** 1 year, 5 months ago  
Had this question in today's exam: 2023-07-26  
upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 8 months ago  
Got this in exam today (5 April 2023)  
upvoted 3 times

🗨️ 👤 **Saluk\_DE** 1 year, 9 months ago  
Question was on exam 2023-03-30  
upvoted 1 times

🗨️ 👤 **proffesormuffin** 1 year, 10 months ago  
400 looks correct to me.

"To estimate the minimum throughput required of a container with manual throughput, find the maximum of:

400 RU/s

Current storage in GB \* 1 RU/s

Highest RU/s ever provisioned on the container / 100"

$5000/100 = 50$  not 500

so  $\text{MAX}(400, 50) = 400$  ?

<https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits>  
upvoted 2 times

🗨️ 👤 **Xardas** 1 year, 5 months ago  
Highest RU/s ever provisioned on the container \* 0.1 = 500  
upvoted 1 times

🗨️ 👤 **Xardas** 1 year, 5 months ago  
This link says formula is  $\text{max} * 0.1$   
<https://learn.microsoft.com/en-us/azure/cosmos-db/how-to-choose-offer#overview-of-provisioned-throughput-types>

But this link says it is  $\text{max} / 100$   
<https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits>

Which is right????  
upvoted 1 times

🗨️ 👤 **stlim83** 9 months, 3 weeks ago  
' $\text{max} / 100$ ', this one is for the manual throughput.  
' $\text{max} / 10$ ', this is for the autoscale throughput.  
upvoted 1 times

🗨️ 👤 **mabdo** 1 year, 11 months ago  
on 2/23  
upvoted 2 times

🗨️ 👤 **Esward** 1 year, 11 months ago  
Given answers are correct.  
upvoted 1 times

🗨️ 👤 **gmishra88** 2 years, 2 months ago  
First option is a question that punishes people who tried to set the RUs and remembered 400 RU as the minimum. If you never tried, you will find some relation with 5000 and 10% is a nice relation.  
upvoted 3 times

🗨️ 👤 **proffesormuffin** 1 year, 10 months ago

it's 1% not 10% so max 400,50 is 400.

<https://learn.microsoft.com/en-us/azure/cosmos-db/concepts-limits>

upvoted 1 times

🗨️ 👤 **fuji36850** 7 months, 4 weeks ago

I think you looked on wrong part of the doc, in 'autoscale' (mentioned in question) it is 10%

upvoted 1 times

🗨️ 👤 **[Removed]** 2 years, 3 months ago

Answer

No, No, Yes, got this in the cosmos dp-420 exam

upvoted 1 times

🗨️ 👤 **Azprep** 2 years, 8 months ago

No,

No,

Yes

upvoted 2 times

## HOTSPOT -

You are developing a web application that makes calls to the Microsoft Graph API. You register the application in the Azure portal and upload a valid X509 certificate.

You create an appsettings.json file containing the certificate name, client identifier for the application, and the tenant identifier of the Azure Active Directory (Azure

AD). You create a method named ReadCertificate to return the X509 certificate by name.

You need to implement code that acquires a token by using the certificate.

How should you complete the code segment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

### Answer Area

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app = .Create(config.ClientId)

    .WithCertificate(certificate)
    .WithAuthority(new Uri(config.Authority))
    .Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient().ExecuteAsync();
```

scopes  
 app  
 config

### Suggested Answer:

#### Answer Area

```
AuthenticationConfig config = AuthenticationConfig.ReadFromJsonFile("appsettings.json");
X509Certificate2 certificate = ReadCertificate(config.CertificateName);
var app = .Create(config.ClientId)

    .WithCertificate(certificate)
    .WithAuthority(new Uri(config.Authority))
    .Build();
string[] scopes = new string[] { $"{config.ApiUrl}.default" };
AuthenticationResult result = await app.AcquireTokenForClient().ExecuteAsync();
```

scopes  
 app  
 config

Box 1: ConfidentialClientApplicationBuilder

Here's the code to instantiate the confidential client application with a client secret: app =

```
ConfidentialClientApplicationBuilder.Create(config.ClientId)
```

```
.WithClientSecret(config.ClientSecret)
```

```
.WithAuthority(new Uri(config.Authority))
```

```
.Build();
```

Box 2: scopes -

After you've constructed a confidential client application, you can acquire a token for the app by calling AcquireTokenForClient, passing the scope, and optionally forcing a refresh of the token.

Sample code: result = await app.AcquireTokenForClient(scopes)

```
.ExecuteAsync();
```

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration> <https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token>

🗄️ 👤 **aradice** Highly Voted 3 years, 6 months ago

correct <https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-app-configuration?tabs=dotnet#instantiate-the-confidential-client-application-with-a-client-certificate>

<https://docs.microsoft.com/en-us/azure/active-directory/develop/scenario-daemon-acquire-token?tabs=dotnet#acquiretokenforclient-api>  
upvoted 55 times

🗄️ 👤 **Esward** 1 year, 11 months ago

Given answers are correct! agreed.

upvoted 2 times

🗄️ 👤 **MK22** Highly Voted 3 years, 4 months ago

Provided answer is correct.

upvoted 13 times

🗄️ 👤 **J\_obando13** Most Recent 1 month, 2 weeks ago

Got this in the exam today! May 15, 2025. Went most voted. 940/1000. 50 of 51 from examtopics. Cases Munson's Pickles and Preserves Farm (7) and Fourth Coffee (3)

upvoted 1 times

🗄️ 👤 **hgbgvfvf** 1 year ago

correct answer

upvoted 1 times

🗄️ 👤 **BaoNguyen2411** 1 year, 5 months ago

got this question on 06-29-2023

upvoted 5 times

🗄️ 👤 **tmuras** 1 year, 10 months ago

This question was on exam on 17/02/23.

upvoted 4 times

🗄️ 👤 **bertwin** 2 years, 1 month ago

Absolute correct

upvoted 1 times

🗄️ 👤 **TheExamMaster2020** 2 years, 1 month ago

Did my exam on 15th November 2022. This question was on it.

upvoted 5 times

🗄️ 👤 **lugospod** 2 years, 11 months ago

Got this one 01/2022. Went with most voted (to avoid writing answers again)

upvoted 9 times