



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



## **CERTIFICATION TEST**

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

## HOTSPOT -

You have an Azure Batch project that processes and converts files and stores the files in Azure storage. You are developing a function to start the batch job.

You add the following parameters to the function.

Parameter name	Description
fileTasks	a list of tasks to be run
jobId	the identifier that must be assigned to the job
outputContainerSasUrl	a storage SAS URL to store successfully converted files
failedContainerSasUrl	a storage SAS URL to store copies of files that failed to convert.

You must ensure that converted files are placed in the container referenced by the outputContainerSasUrl parameter. Files which fail to convert are placed in the container referenced by the failedContainerSasUrl parameter.

You need to ensure the files are correctly processed.

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:

**Suggested Answer:** *Explanation*

Box 1: CreateJob -

Box 2: TaskSuccess -

TaskSuccess: Upload the file(s) only after the task process exits with an exit code of 0.

Incorrect: TaskCompletion: Upload the file(s) after the task process exits, no matter what the exit code was.

Box 3: TaskFailure -

TaskFailure: Upload the file(s) only after the task process exits with a nonzero exit code.

Box 4: OutputFiles -

To specify output files for a task, create a collection of OutputFile objects and assign it to the CloudTask.OutputFiles property when you create the task.

References:

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.batch.protocol.models.outputfileuploadcondition>

<https://docs.microsoft.com/en-us/azure/batch/batch-task-output-files>

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

options

[https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_013\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_013_Image_0001.jpg)

upvoted 36 times

 **Nehagg13** Highly Voted 4 years, 10 months ago

Where are the block/hot area .. Question is incomplete

upvoted 11 times

 **jhonattansolarte** Most Recent 3 years, 10 months ago

<https://user-images.githubusercontent.com/8276103/125325099-2c736100-e306-11eb-8cd9-3baec56235cc.jpeg>

upvoted 1 times

 **chanchal** 4 years, 3 months ago

I cannot see the image.

upvoted 1 times

 **spookyfox** 4 years, 2 months ago



<https://vceguide.com/how-should-you-complete-the-code-segment-25/>

upvoted 2 times

 **Aqauresnight** 4 years, 5 months ago

If your block or answer area is not loaded go to the link at gssicftohbcrsgtjg he gave the image of the answer

upvoted 1 times

  **diligent176** 4 years, 4 months ago



This site is just terrible. Site admins - if you're asking people to pay for access... The site should have working page content and correct answers would be OK too.

upvoted 4 times

  **danishasytarazizan** 4 years, 9 months ago



Why I cannot see the image ?

upvoted 4 times

  **KP\_7082** 4 years, 10 months ago

I cannot see the question for this

upvoted 5 times

  **PraPat** 5 years, 4 months ago

Answer given is correct.

upvoted 7 times

You are writing code to create and run an Azure Batch job.

You have created a pool of compute nodes.

You need to choose the right class and its method to submit a batch job to the Batch service.

Which method should you use?

- A. `JobOperations.EnableJobAsync(String, IEnumerable<BatchClientBehavior>, CancellationToken)`
- B. `JobOperations.CreateJob()`
- C. `CloudJob.Enable(IEnumerable<BatchClientBehavior>)`
- D. `JobOperations.EnableJob(String, IEnumerable<BatchClientBehavior>)`
- E. `CloudJob.CommitAsync(IEnumerable<BatchClientBehavior>, CancellationToken)`

**Suggested Answer: E**

A Batch job is a logical grouping of one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. The app uses the

`BatchClient.JobOperations.CreateJob` method to create a job on your pool.



The Commit method submits the job to the Batch service. Initially the job has no tasks.

```
{
CloudJob job = batchClient.JobOperations.CreateJob();
job.Id = JobId;
job.PoolInformation = new PoolInformation { PoolId = PoolId }; job.Commit();
}
```

...


References:

<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>

  **justanotherusername** Highly Voted 5 years, 4 months ago

E is the answer, `JobOperations.CreateJob()` creates a job that has not been submitted to the Batch service, the question asks for the method that SUBMITS the job to the service.

upvoted 23 times

  **Juanlu** 4 years, 5 months ago

Agree !

upvoted 3 times

  **lywwind** Highly Voted 5 years, 4 months ago

Answer is E

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.batch.cloudjob.commitasync?view=azure-dotnet>

Commits this CloudJob to the Azure Batch service.

upvoted 13 times

  **Larry88** Most Recent 4 years, 8 months ago



tested below code, answer is E

```
try
{
CloudJob job = batchClient.JobOperations.CreateJob();
job.Id = JobId;
job.PoolInformation = new PoolInformation { PoolId = PoolId };
```

```
job.CommitAsync();
```

```
}
```

upvoted 2 times

  **Larry88** 4 years, 8 months ago

tested below code, answer is E

try

```
{  
    CloudPool pool = batchClient.PoolOperations.CreatePool(  
        poolId: PoolId,  
        targetDedicatedComputeNodes: PoolNodeCount,  
        virtualMachineSize: PoolVMSize,  
        virtualMachineConfiguration: vmConfiguration);  
  
    pool.CommitAsync();  
}
```

upvoted 2 times

🗲️ 👤 **navaneeth** 4 years, 9 months ago

B doesn't make sense. It Clears says Submits the job. How come B?

upvoted 1 times

🗲️ 👤 **sathish11** 4 years, 10 months ago

Correct answer is E

upvoted 1 times

🗲️ 👤 **tim291** 4 years, 10 months ago

Yes E is correct

upvoted 1 times

🗲️ 👤 **gssicftohbcrcsgtjg** 5 years, 3 months ago

E is correct

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.batch.cloudjob.commitasync?view=azure-dotnet>

The Commit method submits the job to the Batch service

<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>

upvoted 10 times

🗲️ 👤 **JohnyT** 5 years, 4 months ago

The question is about submitting the job, so answer E.

upvoted 13 times

🗲️ 👤 **Nagarajan** 5 years, 4 months ago

Answer is B and E

upvoted 2 times

🗲️ 👤 **PraPat** 5 years, 4 months ago

Yes, ANS is B. Check <https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>

upvoted 1 times

🗲️ 👤 **AzureGC** 5 years, 4 months ago

Answer # does NOT match the solution description; The correct answer should be "B".,

<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet#review-the-code>

upvoted 5 times

DRAG DROP -

You are developing Azure WebJobs.

You need to recommend a WebJob type for each scenario.

Which WebJob type should you recommend? To answer, drag the appropriate WebJob types to the correct scenarios. Each WebJob 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

WebJob types	Scenario	WebJob type
Triggered	Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance.	
Continuous	Run on a single instance that Azure select for load balancing.	
	Supports remote debugging	

Suggested Answer:

## Answer Area

WebJob types	Scenario	WebJob type
Triggered	Run on all instances that the web app runs on. Optionally restrict the WebJob to a single instance.	Continuous
Continuous	Run on a single instance that Azure select for load balancing.	Triggered
	Supports remote debugging	Continuous

Box 1: Continuous -

Continuous runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.

Box 2: Triggered -

Triggered runs on a single instance that Azure selects for load balancing.

Box 3: Continuous -

Continuous supports remote debugging.

Note:

The following table describes the differences between continuous and triggered WebJobs.

Continuous	Triggered
Starts immediately when the WebJob is created. To keep the job from ending, the program or script typically does its work inside an endless loop. If the job does end, you can restart it.	Starts only when triggered manually or on a schedule.
Runs on all instances that the web app runs on. You can optionally restrict the WebJob to a single instance.	Runs on a single instance that Azure selects for load balancing.
Supports remote debugging.	Doesn't support remote debugging.

References:

<https://docs.microsoft.com/en-us/azure/app-service/web-sites-create-web-jobs>

I got 900 today, my ans: all inst ->

Continuos

Trigger

Continuos

upvoted 13 times

  **Ummara** 4 years, 9 months ago

Congratulation :)

upvoted 1 times

  **Cholan**  4 years, 11 months ago

Right Answer.

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

upvoted 6 times

You are developing a software solution for an autonomous transportation system. The solution uses large data sets and Azure Batch processing to simulate navigation sets for entire fleets of vehicles.

You need to create compute nodes for the solution on Azure Batch.

What should you do?

- A. In the Azure portal, add a Job to a Batch account.
- B. In a .NET method, call the method: `BatchClient.PoolOperations.CreateJob`
- C. In Python, implement the class: `JobAddParameter`
- D. In Azure CLI, run the command: `az batch pool create`
- E. In a .NET method, call the method: `BatchClient.PoolOperations.CreatePool`
- F. In Python, implement the class: `TaskAddParameter`
- G. In the Azure CLI, run the command: `az batch account create`

**Suggested Answer: B**

A Batch job is a logical grouping of one or more tasks. A job includes settings common to the tasks, such as priority and the pool to run tasks on. The app uses the

`BatchClient.JobOperations.CreateJob` method to create a job on your pool.

Note:

Step 1: Create a pool of compute nodes. When you create a pool, you specify the number of compute nodes for the pool, their size, and the operating system.

When each task in your job runs, it's assigned to execute on one of the nodes in your pool.

Step 2: Create a job. A job manages a collection of tasks. You associate each job to a specific pool where that job's tasks will run.

Step 3: Add tasks to the job. Each task runs the application or script that you uploaded to process the data files it downloads from your Storage account. As each task completes, it can upload its output to Azure Storage.

Incorrect Answers:

C, F: To create a Batch pool in Python, the app uses the `PoolAddParameter` class to set the number of nodes, VM size, and a pool configuration.

E: `BatchClient.PoolOperations` does not have a `CreateJob` method.

References:


<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet> <https://docs.microsoft.com/en-us/azure/batch/quick-run-python>

 **jddjsjdsf** Highly Voted 5 years, 6 months ago


I believe D and E are correct. to create compute nodes, you need to have a pool. and D and E can create a pool  
upvoted 19 times

 **James007** Highly Voted 5 years, 3 months ago


B is "incorrect" - there is no `BatchClient.PoolOperations.CreateJob`.  
The correct answer is E - `BatchClient.PoolOperations.CreatePool`.  
upvoted 11 times

 **godsbane** 4 years, 11 months ago

True. Just tested this on Microsoft.Azure.Batch SDK 13.0.0.  
upvoted 2 times

 **Cornholioz** Most Recent 4 years, 5 months ago

Namespace for `BatchClient` class is `Batch`. It does have the `PoolOperations.CreatePool` method.  
<https://csharp.hotexamples.com/examples/-/BatchClient/-/php-batchclient-class-examples.html#0x8180467f32700c3e40f78385dc24d16909fa1afab4b228f42ea4ee64d0e96ab3-16,43>,  
upvoted 2 times

 **Larry88** 4 years, 8 months ago

tested below code, D, E are correct

try

{

CloudPool pool = batchClient.PoolOperations.CreatePool(







```
poolId: PoolId,  
targetDedicatedComputeNodes: PoolNodeCount,  
virtualMachineSize: PoolVMSize,  
virtualMachineConfiguration: vmConfiguration);
```



```
pool.CommitAsync();  
}
```



```
az batch account login \  
--resource-group batchRG \  
--name batchwerviceacc  
--shared-key-auth
```

```
az batch pool create  
upvoted 6 times
```

  **Juanlu** 4 years, 2 months ago  
I Agree !  
upvoted 1 times

  **sathish11** 4 years, 10 months ago  
I think correct answer is E.  
<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>  
upvoted 2 times



  **Nick\_1** 4 years, 10 months ago  
"You need to create compute nodes for the solution on Azure Batch"  
do we have a batch account?  
upvoted 1 times



  **niloySubs** 4 years, 10 months ago  
We have confusions over B, D, E  
B >> Clearly Wrong Answer. As PoolOperations class do not have any method like "CreateJob"



D >> az batch pool create. This creates a Batch Pool to an Account !!



E >> PoolOperations.CreatePool() this one creates a Batch Pool too. BUT that is unbound and has NO relationship to any batchService\Account. We must use commitAsyncl() to add the pool to a Batch Account.

based on these, D seems to the more-correct answer.  
upvoted 3 times

  **Suchismita** 4 years, 11 months ago  
Both D & E are correct , syntax's will create a pool of compute node in CLI and .net .  
upvoted 2 times

  **AK89** 5 years ago  
The Correct Answer is D and E as here you need to create only the pool and not worry of rest of the logic. The link has the official code snippets for your reference <https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>  
upvoted 1 times

  **Nanaboso2** 5 years ago  
You state in your own answer "E: BatchClient.PoolOperations does not have a CreateJob method." Unfortunately "BatchClient.PoolOperations.CreateJob" is associated with B. Please update your solution to mark E as the correct answer  
upvoted 2 times

  **barubs** 5 years ago  
In my opinion, the correct answer is D. Why? "...on Azure Batch" means that correct answer should create pool on Azure. Ans. E, is only partially good solution because it does not contain second step: CommitAsync which is required to submit this command to Azure.  
upvoted 7 times

  **lloves73g** 5 years ago

Answer is D:

<https://docs.microsoft.com/en-us/cli/azure/batch/pool?view=azure-cli-latest#az-batch-pool-create>

Why E is not correct?

Because :

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.batch.pooloperations?view=azure-dotnet>



There is a `CreateComputeNodeUser(String, String) ...`

upvoted 3 times

  **markrobinson** 5 years, 1 month ago


Answer is D & E

upvoted 3 times

  **heero** 5 years, 2 months ago

I think D and E are correct

upvoted 4 times

  **KK\_uniq** 5 years, 3 months ago

In a .NET method, call the method: `BatchClient.PoolOperations.CreatePool`

This is the one which will get compute nodes

upvoted 4 times

  **lywwind** 5 years, 4 months ago

E

<https://docs.microsoft.com/en-us/azure/batch/quick-run-dotnet>

Create a pool of compute nodes

To create a Batch pool, the app uses the `BatchClient.PoolOperations.CreatePool` method to set the number of nodes, VM size, and a pool configuration.



upvoted 2 times

  **AzureGC** 5 years, 4 months ago

E only, as the scenario is describing a software solution, that runs autonomous;

Still need to create a batch job, at some point, but the question is ONLY asking about creating the "pool";

upvoted 5 times

  **cbn** 4 years, 3 months ago

This is a good point. It has to be a software solution I guess and not a manual step.

upvoted 1 times

DRAG DROP -

You are deploying an Azure Kubernetes Services (AKS) cluster that will use multiple containers.

You need to create the cluster and verify that the services for the containers are configured correctly and available.

Which four commands should you use to develop the solution? To answer, move the appropriate command segments from the list of command segments to the answer area and arrange them in the correct order.

Select and Place:

### Command segments

### Answer Area

az aks get-credentials

az appservice plan create

az aks create

az group create

kubectl apply



Suggested Answer:

### Command segments

### Answer Area

az aks get-credentials

az appservice plan create

az aks create

az group create

kubectl apply

az group create

az aks create

kubectl apply

az aks get-credentials



Step 1: az group create -

Create a resource group with the az group create command. An Azure resource group is a logical group in which Azure resources are deployed and managed.

Example: The following example creates a resource group named myAKSCluster in the eastus location. az group create --name myAKSCluster --location eastus

Step 2 : az aks create -

Use the az aks create command to create an AKS cluster.

Step 3: kubectl apply -

To deploy your application, use the kubectl apply command. This command parses the manifest file and creates the defined Kubernetes objects.

Step 4: az aks get-credentials -

Configure it with the credentials for the new AKS cluster. Example: az aks get-credentials --name aks-cluster --resource-group aks-resource-group

References:

<https://docs.bitnami.com/azure/get-started-aks/>

az group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 77 times



  **aes**  5 years, 6 months ago

az aks get-credentials should be executed before any kubectl command in order to get access credentials for the created cluster  
upvoted 48 times


  **NhanTran1997**  1 year, 1 month ago

It should be  
az group create  
az aks create  
az aks get-credentials  
kubectl apply

Ref: <https://learn.microsoft.com/en-us/azure/aks/tutorial-kubernetes-deploy-cluster?tabs=azure-cli>  
upvoted 2 times



  **hergoi** 4 years, 2 months ago

Please correct answer to:  
az-group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 2 times

  **MaTaO** 4 years, 6 months ago

az-group create  
az aks create  
az aks get-credentials  
kubectl apply

confirm, I run a solution on aks since start of 2019. need to get credential before apply\*  
upvoted 1 times

  **hertino** 4 years, 9 months ago


Today I got 900, I had this question, my answer:  
az-group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 8 times

  **navaneeth** 4 years, 9 months ago

admin change the answer to  
az group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 6 times

  **RajeshSinha** 4 years, 10 months ago

Correct answer is:  
az group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 6 times

  **sathish11** 4 years, 10 months ago

Correct answer is

az group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 5 times

  **dollarpo7** 5 years ago

az group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 5 times

  **vova12345** 5 years ago

incorrect answer, correct:  
az group create  
az aks create  
az aks get-credentials  
kubectl apply  
upvoted 14 times

## DRAG DROP -

You are preparing to deploy a medical records application to an Azure virtual machine (VM). The application will be deployed by using a VHD produced by an on-premises build server.

You need to ensure that both the application and related data are encrypted during and after deployment to Azure.

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:

**Suggested Answer:** *Explanation*

Step 1: Encrypt the on-premises VHD by using BitLocker without a TPM. Upload the VM to Azure Storage

Step 2: Run the Azure PowerShell command Set-AzureRMVMOSDisk

To use an existing disk instead of creating a new disk you can use the Set-AzureRMVMOSDisk command.

Example:

```
$osDiskName = $vmname+'_osDisk'
```

```
$osDiskCaching = 'ReadWrite'
```

```
$osDiskVhdUri = "https://$storage.blob.core.windows.net/vhds/" + $vmname + "_os.vhd"
```

```
$vm = Set-AzureRmVMOSDisk -VM $vm -VhdUri $osDiskVhdUri -name $osDiskName -Create
```

Step 3: Run the Azure PowerShell command Set-AzureRmVMDiskEncryptionExtension

Use the Set-AzVMDiskEncryptionExtension cmdlet to enable encryption on a running IaaS virtual machine in Azure.

Incorrect:

Not TPM: BitLocker can work with or without a TPM. A TPM is a tamper resistant security chip on the system board that will hold the keys for encryption and check the integrity of the boot sequence and allows the most secure BitLocker implementation. A VM does not have a TPM.

References:

<https://www.itprotoday.com/iaaspaas/use-existing-vhd-azurerem-vm>

 **PraPat**  5 years, 4 months ago

Check <https://vceguide.com/which-three-actions-should-you-perform-in-sequence-606/> for options  
upvoted 21 times

 **R2C** 5 years, 1 month ago

The link does not work.  
upvoted 3 times

 **Raevox** 4 years, 11 months ago

The link worked fine for me. What is nice is the web page provides a graphic of the problem that is missing in this examtopics.com question.  
upvoted 4 times

 **Not\_Sure\_Who\_Am\_I**  4 years, 11 months ago

Options: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_019\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_019_Image_0001.jpg)  
upvoted 10 times

 **jhonattansolarte**  3 years, 10 months ago

RESPONSE:

<https://user-images.githubusercontent.com/8276103/125327598-e8359000-e308-11eb-8ee2-3c4c02a5a276.jpeg>  
upvoted 1 times


 **AnonymousJhb** 4 years, 5 months ago

The answer is incorrect.

as per the MS documentation, Set-AzureRmVMDiskEncryptionExtension Enables encryption on a running IaaS virtual machine in Azure.

Thus, we need to first provision a New-AzVm in step 3 and then step 4 Set-AzureRmVMDiskEncryptionExtension

<https://docs.microsoft.com/en-us/powershell/module/azurerem.compute/set-azurermvmdiskencryptionextension?view=azurermps-6.13.0>  
upvoted 1 times

 **altafpatel1984** 3 years, 6 months ago

But question asks for only 3 steps, not 4 steps  
upvoted 1 times

 **Cornholioz** 4 years, 5 months ago



But the question says "encrypted during and after deployment to Azure". It does ask for an encryption step "After Deployment to Azure". Poorly framed question, but I'm guessing it is looking for the step to encrypt AFTER deployment to Azure which is the Set-AzVMDiskEncryptionExtension I'm inclined to going with given answer.

upvoted 1 times

  **NajamKhan** 4 years, 6 months ago

Options: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_019\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_019_Image_0001.jpg)

upvoted 1 times



  **Dirk** 4 years, 9 months ago

I'm not so sure about the TPM.

It should be possible as well with a VM:

<https://www.prajwaldesai.com/enable-trusted-platform-module-on-virtual-machine/>

upvoted 1 times

  **agoncal** 4 years, 11 months ago

Last option should be New=AzureRmVM, encryption can be set in Set-AzureRMVMOsDisk cmd:

"Encrypt VMs with pre-encrypted VHDs with Azure PowerShell

You can enable disk encryption on your encrypted VHD by using the PowerShell cmdlet Set-AzVMOSDisk. The example below gives you some common parameters.

PowerShell

```
$VirtualMachine = New-AzVMConfig -VMName "MySecureVM" -VMSize "Standard_A1"
```

```
"
```

```
$VirtualMachine = Set-AzVMOSDisk -VM $VirtualMachine -Name "SecureOSDisk" -VhdUri "os.vhd" Caching ReadWrite -Windows -CreateOption "Attach" -DiskEncryptionKeyUrl "https://mytestvault.vault.azure.net/secrets/Test1/514ceb769c984379a7e0230bddaaaaaa" -DiskEncryptionKeyVaultId "/subscriptions/00000000-0000-0000-0000-000000000000/resourceGroups/myKVresourcegroup/providers/Microsoft.KeyVault/vaults/mytestvault"
```

```
New-AzVM -VM $VirtualMachine -ResourceGroupName "MyVirtualMachineResourceGroup"
```

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disk-encryption-windows-aad>

upvoted 5 times

  **rroongcheng** 4 years, 10 months ago

Yes, It should be:

1. get VHD without TPM

(\$VirtualMachine = New-AzureRmVMConfig) not include in option

2.Set-AzVMOSDisk -VM \$VirtualMachine -VhdUri "os.vhd"

3.New-AzureRmVM -VM \$VirtualMachine

<https://docs.microsoft.com/en-us/powershell/module/AzureRM.Compute/Set-AzureRmVMOSDisk?view=azurermps-6.13.0#examples>

upvoted 9 times

  **tsimbah** 5 years, 1 month ago

I ok ok

I

upvoted 1 times

  **hmirko** 5 years, 2 months ago



Is last step New-AzureRmVM or Set-AzureRmVMDiskEncryptionExtension?

upvoted 3 times

  **VMCoder** 5 years, 2 months ago

I think it should be New-AzureRmVM as well. We have already encrypted the disk in first step.

upvoted 6 times

  **LTiwana** 4 years, 10 months ago

looks correct to me as mentioned in this document:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disk-encryption-windows#new-iaas-vms-created-from-customer-encrypted-vhd-and-encryption-keys>

upvoted 1 times

🗨️ 👤 **Cornholioz** 4 years, 5 months ago

New-AzureRmVM is correct.

NOT Set-AzureRmVMDiskEncryptionExtension because it enables encryption on a "running" IaaS virtual machine in Azure. In this case the VM is not created yet. We have only prepped the OS Disk.

upvoted 1 times

🗨️ 👤 **Cornholioz** 4 years, 5 months ago

Rethinking: But the question says "encrypted during and after deployment to Azure". It does ask for an encryption step "After

Deployment to Azure". Poorly framed question, but I'm guessing it is looking for the step to encrypt AFTER deployment to Azure which is the Set-AzVMDiskEncryptionExtension

I'm inclined to going with given answer.

upvoted 1 times

🗨️ 👤 **wlfjck** 5 years, 3 months ago

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/disk-encryption-sample-scripts>

upvoted 1 times



## DRAG DROP -

You plan to create a Docker image that runs as 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 setupScript.ps1 when the container is built.
- ⇒ Run ContosoApp.dll when the container starts.

The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.

Which four 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:

**Suggested Answer:** *Explanation*

Step 1: WORKDIR /apps/ContosoApp

Step 2: COPY ./-

The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored.

Step 3: EXPOSE ./ContosoApp/ /app/ContosoApp

Step 4: CMD powershell ./setupScript.ps1

ENTRYPOINT ["dotnet", "ContosoApp.dll"]

You need to create a Dockerfile document that meets the following requirements:

- ⇒ Call setupScript.ps1 when the container is built.
- ⇒ Run ContosoApp.dll when the container starts.

References:

<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image>

 **sho1**  5 years, 3 months ago

FROM

WORKDIR

COPY

RUN


CMD

upvoted 41 times

 **riley5** 5 years, 1 month ago


Why not CMD + ENTRYPOINT?

upvoted 2 times

 **BakePVL** 4 years, 11 months ago


CMD doesn't run powershell command when container is built. It only runs command when container starts. <https://goinbigdata.com/docker-run-vs-cmd-vs-entrypoint/>

upvoted 2 times

 **zenth0r** 4 years, 11 months ago

CMD and ENTRYPOINT differ only slightly, and are run in the runtime. The requirement is to run the powershell script during build, hence RUN.

upvoted 1 times

 **riley5** 5 years, 1 month ago


Run + cmd executes build time instructions, but doesn't run in the container

upvoted 2 times

 **tejasmehta8819** 4 years, 4 months ago

As per question, 4 commands to be used not 5

upvoted 2 times

 **cbn** 4 years, 3 months ago

Agree. A good reading on selecting the last option:

<https://goinbigdata.com/docker-run-vs-cmd-vs-entrypoint/>

upvoted 1 times

🗨️ 👤 **gssicftohbcrcsgtjg** Highly Voted 5 years, 3 months ago

Options

<https://www.examtopycs.com/assets/media/exam-media/02758/0036700001.jpg>

upvoted 27 times

🗨️ 👤 **jhonattansolarte** Most Recent 3 years, 10 months ago

\*\*\*\*

RESPONSE:

<https://user-images.githubusercontent.com/8276103/125328006-6abe4f80-e309-11eb-8f4d-ff78414aff3a.jpeg>

upvoted 1 times

🗨️ 👤 **simonverma** 4 years, 1 month ago

STEP 1: FROM

STEP 2: WORKDIR

STEP 3: COPY

STEP 4: RUN POWERSHELL CMD

Image Url: <https://www.examtopycs.com/discussions/microsoft/view/13131-exam-az-300-topic-3-question-4-discussion/>

upvoted 1 times

🗨️ 👤 **bhushan\_786** 4 years, 6 months ago

The answer provided by examtopycs is totally different than the discussions here, can someone provide correct answer please?

upvoted 1 times

🗨️ 👤 **leshuva** 4 years, 5 months ago

FROM base\_image

WORKDIR tocreatedirectoryinthecontainer

COPY thecontentfromexistingcodedirectorytocontainerimage

CMD

upvoted 2 times

🗨️ 👤 **mszydel** 4 years, 8 months ago

There are four commands to choose, most of you suggest 5 commands ...

upvoted 3 times

🗨️ 👤 **Subhijith** 4 years, 9 months ago

FROM base\_image

WORKDIR tocreatedirectoryinthecontainer

COPY thecontentfromexistingcodedirectorytocontainerimage

CMD

ENTRYPOINT

upvoted 4 times

🗨️ 👤 **hertino** 4 years, 9 months ago

Today I got 900, I had this question, my answer:

WORKDIR /apps/ContosoApp

COPY ./-

EXPOSE ./ContosApp/ /app/ContosoApp

CMD powershell ./setupScript.ps1

ENTRY

upvoted 2 times

🗨️ 👤 **SilNilanjan** 4 years, 9 months ago

Don't think so... there is no base-image, plus the EXPOSE command should have this syntax: EXPOSE <port> [<port>/<protocol>...] but it is incorrect here

upvoted 1 times

🗨️ 👤 **bhushan\_786** 4 years, 4 months ago

This could be one of your wrong answer :-)

upvoted 3 times

🗨️ 👤 **Steve\_az** 4 years, 9 months ago

<https://www.examtopycs.com/discussions/microsoft/view/13131-exam-az-300-topic-3-question-4-discussion/>

upvoted 2 times

🗨️ 👤 **brajen** 4 years, 10 months ago

The last command will be used to run the application when the container starts up. We have to use the ENTRYPOINT command for this.

So the correct answer is-

FROM > WORKDIR > COPY > CMD/ENTRYPOINT

upvoted 1 times

🗨️ 👤 **Not\_Sure\_Who\_Am\_I** 4 years, 11 months ago

Options: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-300-date-01-06-2019-00001\\_Page\\_304\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-300-date-01-06-2019-00001_Page_304_Image_0001.jpg)

upvoted 4 times

🗨️ 👤 **AK89** 5 years ago

The correct answer is :

FROM - this is always starting point of a Dockerfile

WORKDIR

COPY

CMD

ENTRYPOINT

upvoted 3 times

🗨️ 👤 **lau13** 5 years ago

This can NOT be the Entrypoint, because the question stated that it must run the setupScript.ps1 when container is built, to do that, it must be the option with RUN

CMD and Entrypoint are similar, they both execute when container starts. The only difference is CMD will be ignored if you execute the container with a command specified; however Entrypoint won't be ignored.

upvoted 3 times

🗨️ 👤 **lloves73g** 5 years ago

FROM WORKDIR COPY RUN CMD

STEP 1: FROM microsoft/ aspnetcore:2.0

STEP 2: WORKDIR /apps/ContosoAPP

STEP 3: COPY . .

STEP 4: RUN powershell setupScript.ps1

CMD ["dotnet", " ContosoApp.dll"]

upvoted 22 times

🗨️ 👤 **shuakwe** 5 years ago

It's asking for four Docker commands, not to choose four options.

"The Docker document must be created in the same folder where ContosoApp.dll and setupScript.ps1 are stored" -> This means WORKDIR is not needed.

Solution:

FROM

COPY

RUN

CMD

upvoted 1 times

🗨️ 👤 **Mvii** 5 years ago

CMD and ENTRYPOINT can't be used together as they do the same thing.

FROM

WORKDIR

COPY

RUN

CMD

upvoted 3 times

🗨️ 👤 **riley5** 5 years, 1 month ago

Correct?

FROM



WORKDIR

COPY

RUN

CMD

upvoted 4 times

  **riley5** 5 years, 1 month ago

correction: FROM > WORKDIR > COPY > CMD/ENTRYPOINT

upvoted 8 times

DRAG DROP -

You are creating a script that will run a large workload on an Azure Batch pool. Resources will be reused and do not need to be cleaned up after use.

You have the following parameters:

Parameter name	Description
\$script	the script that will run across the batch pool
\$image	the image that pool worker processes will use
\$sku	the node agent SKU Id
\$numberOfJobs	the number of jobs to run

You need to write an Azure CLI script that will create the jobs, tasks, and the pool.

In which order should you arrange the commands to develop the solution? To answer, move the appropriate commands from the list of command segments to the answer area and arrange them in the correct order.

Select and Place:

**Suggested Answer:** *Explanation*

Step 1: az batch pool create -

```
# Create a new Linux pool with a virtual machine configuration. az batch pool create \
--id mypool \
--vm-size Standard_A1 \
--target-dedicated 2 \
--image canonical:ubuntuserver:16.04-LTS \
--node-agent-sku-id "batch.node.ubuntu 16.04"
```

Step 2: az batch job create -

```
# Create a new job to encapsulate the tasks that are added.
az batch job create \
--id myjob \
--pool-id mypool
```


Step 3: az batch task create -

```
# Add tasks to the job. Here the task is a basic shell command. az batch task create \
--job-id myjob \
--task-id task1 \
--command-line "/bin/bash -c 'printenv AZ_BATCH_TASK_WORKING_DIR'"
```

Step 4: for i in {1..\$numberOfJobs} do

References:

<https://docs.microsoft.com/bs-latn-ba/azure/batch/scripts/batch-cli-sample-run-job>

 **Dann1112** Highly Voted 5 years, 3 months ago

Should be:

Create Pool

Create Job

For loop

Create Task

Reference: <https://docs.microsoft.com/bs-latn-ba/azure/batch/quick-create-cli#create-tasks>

upvoted 66 times

 **Daltonic75** 5 years, 2 months ago

Agree. "Create task" uses "\$i" concatenated to the task ID value, where "i" is the loop index.

upvoted 9 times


 **Juanlu** 4 years, 5 months ago

Agree !

upvoted 1 times

 **PraPat** Highly Voted 5 years, 4 months ago

Check <https://vceguide.com/which-order-should-you-arrange-the-commands-to-develop-the-solution/> for options  
upvoted 17 times

  **DanielNadjbabi** 4 years, 5 months ago

<https://vceguide.com/which-order-should-you-arrange-the-commands-to-develop-the-solution/>  
upvoted 2 times

  **jhonattansolarte** Most Recent 3 years, 10 months ago

\*\*\* RESPONSE (Image)



<https://user-images.githubusercontent.com/8276103/125328926-7fe7ae00-e30a-11eb-8bb9-55e5c3932e0b.jpeg>  
upvoted 1 times

  **Deputy\_Cartman** 4 years, 1 month ago

Another incorrect answer on this website. In other news, water remains wet.

1. az batch pool create
2. az batch job create
3. for i in {1..4}
4. az batch task create

<https://docs.microsoft.com/en-us/azure/batch/quick-create-cli>  
Last updated: 13 August 2020  
<https://docs.microsoft.com/en-us/azure/batch/quick-create-cli>  
upvoted 1 times

  **ss22** 4 years, 3 months ago

<https://docs.microsoft.com/bs-latn-ba/azure/batch/quick-create-cli#create-tasks>  
upvoted 1 times

  **SilNilanjan** 4 years, 9 months ago

The order given in the question should be the correct order, there is nothing to alter  
upvoted 4 times

  **triptimandal01** 4 years, 10 months ago

I agree for loop should be #3 and var i is refereed in create task commands. So which means the order in question is right order..no need to change.  
upvoted 1 times

  **AK89** 5 years ago

Should be:

Create Pool

Create Job

For loop

Create Task

upvoted 7 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($"Queue Expiration Time: {myQueueItem.ExpirationTime}");
        tableBindings.Add(JsonConvert.DeserializeObject<Order>(myQueueItem.AsString));
    }
    [FunctionName("ProcessOrders-Poison")]
    public static void ProcessFailedOrders([QueueTrigger("incoming-orders-poison")]CloudQueueMessage myQueueItem, TraceWriter log)
    {
        logError($"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.	<input type="radio"/>	<input type="radio"/>
When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.	<input type="radio"/>	<input type="radio"/>
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.	<input type="radio"/>	<input type="radio"/>
The ProcessOrders function will output the order to an Orders table in Azure Table Storage.	<input type="radio"/>	<input type="radio"/>

Suggested Answer:

### Answer Area

	Yes	No
The code will log the time that the order was processed from the queue.	<input type="radio"/>	<input checked="" type="radio"/>
When the ProcessOrders function fails, the function will retry up to five times for a given order, including the first try.	<input checked="" type="radio"/>	<input type="radio"/>
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.	<input checked="" type="radio"/>	<input type="radio"/>
The ProcessOrders function will output the order to an Orders table in Azure Table Storage.	<input checked="" type="radio"/>	<input type="radio"/>

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 -



References:

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue>

  **Khang** Highly Voted 5 years, 4 months ago

Answer is correct

upvoted 19 times

  **Juanlu** 4 years, 5 months ago


Agree !

upvoted 1 times

  **RahulKate** Most Recent 4 years, 4 months ago

5 times retry - <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-storage-queue-trigger?tabs=csharp#poison-messages>

upvoted 3 times

  **ijN00M4** 4 years, 10 months ago

Doesn't `log.Info` include a timestamp?

Meaning `log.info(\$"Processing Order: {myQueueItem.Id}")` should show as

`<TIMESTAMP> [Information] Processing Order: <OrderId>`, and thus the processing time is the timestamp

upvoted 2 times

  **Tealon** 4 years, 5 months ago

Yes, but that is not done by the code, but by the Azure function that stores the logs.

upvoted 2 times

  **sathish11** 4 years, 10 months ago

Answer is correct

upvoted 3 times

  **NKnab** 4 years, 11 months ago

By default, blob, queue and service triggers retry requests up to five times. After the fifth retry, both the Azure Queue storage and Azure Service Bus triggers write a message to a poison queue.

upvoted 3 times

  **ashtwentyfour** 4 years, 11 months ago

Regarding retries - <https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages>

upvoted 2 times

  **sjsouvik** 5 years ago

where is maxDequeueCount ( The number of times to try processing a message before moving it to the poison queue. Default value is 5) part in the code ? so I believe it should be marked as "NO".

upvoted 3 times

  **Grefz** 4 years, 9 months ago

The trigger is set as default to retry up to five times.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-bindings-error-pages>

upvoted 5 times

  **silentsouls** 4 years, 3 months ago

but that states: if using Azure Service Bus, by default queues have a message delivery count of 10.

upvoted 1 times



  **vova12345** 5 years ago



Incorrect answer, once again ;(  
myQueueItem.AsString contains the time of adding the message to the queue  
upvoted 2 times

  **[Removed]** 5 years ago

That is the insertion time not processing time.  
upvoted 5 times

  **oxaytol** 4 years, 11 months ago

Answer is correct as the time of addition is not the processing time. You can add an item to a queue but processing could happen some time later.  
upvoted 1 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.

References:

<https://docs.microsoft.com/mt-mt/azure/app-service/containers/quickstart-docker-go?view=sql-server-ver15>

**hertino** Highly Voted 4 years, 9 months ago

Today I got 900, I had this question, my answer:

az group create -

az appservice plan create

az webapp create -

upvoted 22 times

🗨️ 👤 **chintan4190** Highly Voted 🏆 5 years, 1 month ago

given answer is correct. In question it is mentioned that None of your current resource groups reside in a location that supports Linux. --> so we have to create RG first

upvoted 17 times

🗨️ 👤 **MaheshBeeravelli** 4 years, 10 months ago

Is there any restriction that app service should be in the same location as in Resource Group ?

upvoted 1 times

🗨️ 👤 **funfun** Most Recent 🔔 4 years, 4 months ago

why not just create a new app service plan in the same resource group but the plan in different region, then create web app? In such case, only 2 steps are needed without creating a new resource group

upvoted 1 times

🗨️ 👤 **SidSadhu** 4 years, 9 months ago

CORRECT

upvoted 1 times

🗨️ 👤 **Jack2016** 4 years, 9 months ago

Answer is correct.

If the resource group does not support linux, new group should be created.

```
{"code": "InvalidTemplateDeployment", "message": "The template deployment 'Microsoft.Web-ASP-Portal-' is not valid according to the validation procedure. See inner errors for details.", "details": [{"code": "ValidationForResourceFailed", "message": "Validation failed for a resource. Check 'Error.Details[0]' for more information.", "details": [{"code": "LinuxWorkersNotAllowedInResourceGroup", "message": "Linux workers are not available in resource group ResoureIndia. Use this link to learn more https://go.microsoft.com/fwlink/?linkid=831180"}]}}
```

upvoted 1 times

🗨️ 👤 **mattos** 5 years, 2 months ago

Answer should be to create app service plan, create webapp, and update webapp. No need to update or create resource groups, since the question mentions there are existing RGs and that number of RGs should be minimized. Resources can be created in any location, regardless of the RG location.

upvoted 1 times

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

makes sense, but what would you use the update webapp for?

upvoted 1 times

🗨️ 👤 **kl337n** 5 years, 2 months ago

You can't mix Linux and Windows app service plans in one resource group, therefore you need to create new RG.

<https://docs.microsoft.com/pl-pl/azure/app-service/containers/app-service-linux-intro>

"Based on a current limitation, for the same resource group you cannot mix Windows and Linux apps in the same region."

upvoted 21 times

🗨️ 👤 **PT29** 4 years, 9 months ago

It says "you cannot mix Windows and Linux apps in the same region." so it can be present in same RG but you will have to create ASP in diff regions for widows and linux. This works and I have it working for me.

upvoted 4 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	Answer Area	
	Requirement	Feature
Run Command	Firewall configuration	
Serial console	Supporting services script	
Hybrid Runbook Worker		
Custom Script Extension		

## Suggested Answer:

Features	Answer Area	
	Requirement	Feature
	Firewall configuration	Run Command
Serial console	Supporting services script	Hybrid Runbook Worker
Custom Script Extension		

## References:

<https://docs.microsoft.com/en-us/azure/automation/automation-hybrid-runbook-worker> <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-command>

 **TechieBloke** Highly Voted 5 years, 4 months ago

firewall: run-command

script: custom script extensions

This capability is useful in all scenarios where you want to run a script within a virtual machine. It's one of the only ways to troubleshoot and remediate a virtual machine that doesn't have the RDP or SSH port open because of improper network or administrative user configuration.



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 35 times

  **divz\_999** 5 years, 4 months ago

Is this correct?



upvoted 1 times

  **Tee48** 5 years, 1 month ago

Yes, It is.



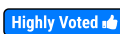
<https://docs.microsoft.com/en-us/azure/virtual-machines/extensions/custom-script-windows>

upvoted 1 times

  **Juanlu** 4 years, 6 months ago

Agree !

upvoted 1 times

  **AK89**  5 years ago

Firewall : Run command

2nd point : Custom Script Extension as the scripts are stored in an Azure Storage Account. "Hybrid Runbook Worker" can be the case only if the scripts are stored in Azure Automation Account which is not the case here.

upvoted 11 times

  **AAAAAAAQQQ** 4 years ago

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm>

upvoted 1 times

  **meoukg**  3 years, 2 months ago

Got it on 03/2022, I chose as Firewall configuration = Run command ; Supporting service scripts = Hybrid Runbook Worker

upvoted 1 times

  **bugimachi** 4 years, 4 months ago

The question asks "WHEN a VM is PROVISIONED the firewall must be AUTOMATICALLY configured". But "Run Command" requires in manual invocation (either via CLI or Portal), doesn't it? So does it really satisfy the requirement of automatic firewall configuration?

I'd personally include it into a script and run it using the Custom Script Extension. What do you think?

upvoted 8 times

  **JulienYork** 4 years, 4 months ago



Same here

upvoted 1 times

  **dluk** 4 years, 4 months ago

But doesn't Custom Script Extension need an access to Azure resources? In the question it states that you need to configure firewall before accessing any Azure resources

upvoted 1 times

  **hertino** 4 years, 9 months ago

Today I got 900, I had this question, my answer:

Firewall configuration Run command



Supporting service Custom script extension

upvoted 10 times

  **XtremeGraff** 4 years, 1 month ago

How many study cases you had ?

upvoted 1 times

  **sathish11** 4 years, 10 months ago

1. Run Command



2. Customer Script Extension

upvoted 2 times

  **Nanaboso2** 5 years ago

From reading <https://docs.microsoft.com/en-us/azure/virtual-machines/windows/run-scripts-in-vm> I think "Custom Script Extension" as well as "Hybrid Runbook Worker" would work for snd. I would appreciate a clarification why "Hybrid Runbook Worker" is preferred a lot.

upvoted 2 times

  **wlfjck** 5 years, 3 months ago

1. Run Command
2. Customer Script Extension

upvoted 9 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](http://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 command from the list of commands to the answer area and arrange them in the correct order.

Select and Place:

**Suggested Answer:** *Explanation*

Step 1: `#bin/bash` -

The `appName` is used when the `webapp-name` is created in step 2.

Step 2: `az webapp config hostname add`

The `webapp-name` is used when the webapp is created in step 3.

Step 3: `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 : `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.

In the Cloud Shell, follow the `az webapp create` command with `az webapp config container set`.

References:

<https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image>

  **AK89**  5 years ago

Step 1: `#bin/bash` -

Step 2 -`az webapp create`

Step 3 -`az webapp config container set`

Step 4- `az webapp config hostname`

The custom domain mapping is done after the app is in place running. So it will be the last step in the process.

Link : <https://docs.microsoft.com/en-us/azure/app-service/containers/tutorial-custom-docker-image>

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web/tutorial-custom-domain?toc=/azure/app-service/containers/toc.json>

upvoted 27 times

  **Avadhut01** 4 years, 11 months ago




Its correct.

upvoted 1 times

  **azurelearner666** 4 years, 11 months ago

Agree, it seems the most correct response. But it seems that in all the exam pdf's it is wrongly evaluated... guess we will see..

upvoted 1 times

  **andreiyfrag**  5 years, 2 months ago

Options: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_050\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_050_Image_0001.jpg)

upvoted 24 times

  **jhonattansolarte**  3 years, 10 months ago

Response (Image)

<https://user-images.githubusercontent.com/8276103/125330387-50d23c00-e30c-11eb-823a-9723728751ad.png>

upvoted 2 times

  **Cornholioz** 4 years, 5 months ago

Looked around quite a bit. Haven't found anything conclusive about the order of the last two steps. I'm going with `config container set` and then `hostname add`.

1. `#bin/bash`

2. `az webapp create`

3. az webapp config container set


4. az webapp config hostname

Samples I'm pasting are not conclusive, but might help someone who wants to code this:

<https://github.com/Azure-Samples/azure-cli-samples/blob/master/app-service/deploy-linux-docker/deploy-linux-docker.sh>

[https://azx.ms/details/webapp\\_container](https://azx.ms/details/webapp_container)

upvoted 3 times

  **Nitin121** 4 years, 7 months ago

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-linux-docker-aspnetcore?toc=/cli/azure/toc.json>

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-configure-custom-domain>

upvoted 2 times

  **Nitin121** 4 years, 7 months ago

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-configure-custom-domain>

correct order for the option

1) Bash

2) Web app create

3) hostname 4) Container

upvoted 2 times

  **Cornholioz** 4 years, 5 months ago

The link is not for container. The ambiguity is on the last 2 steps. Should the Container be set before or after configuring the Host name. I haven't tested it with a PS myself but for now, I'm going with Container Set config before Host name config.

upvoted 2 times

  **NaliniTomar** 4 years, 9 months ago


<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-configure-custom-domain>

upvoted 1 times

  **Steve\_az** 4 years, 9 months ago

Refer this link for answer options - <https://www.examttopics.com/assets/media/exam-media/02758/0036400001.jpg>

upvoted 1 times

  **Iloves73g** 5 years ago

Step 1: #bin/bash -

Step 2 -az webapp create

Step 3 -az webapp config container set

Step 4- az webapp config hostname

if without az webapp config container set, webapp is not created, here config hostname has not place to do...

upvoted 7 times

  **[Removed]** 5 years, 1 month ago

Shouldn't the order be:

Step 1: #bin/bash -

Step 2 -az webapp create

Step 3 -az webapp config container set

Step 4- az webapp config hostname

So that there is minimal downtime for anyone accessing the app via the domain name?

upvoted 10 times

  **tiger123** 5 years, 1 month ago

Step 1. bin/bash

Step 2. webapp create

Step 3. hostname add

Step 4. config container set

upvoted 3 times

  **tiger123** 5 years, 1 month ago

Step1. #bin/bash

Step2. az webapp create

Step 3. az webapp create

Step 4. az webapp config container set



upvoted 1 times

🗨️ 👤 **heero** 5 years, 2 months ago

-az webapp create  
az webapp confing container set  
#bin/bash -  
az webapp config hostname  
upvoted 1 times

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

wrong  
-az webapp create need \$appName  
upvoted 1 times

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

someone please post the options.  
upvoted 3 times

🗨️ 👤 **heero** 5 years, 3 months ago

Step 1: #bin/bash -  
Step 2 -az webapp create  
Step 3 -az webapp confing container set  
Step 4- az webapp config hostname  
upvoted 19 times

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

why? az webapp config container set stands for docker container update. you would want to configure app first and update container later.  
upvoted 2 times

🗨️ 👤 **heero** 5 years, 3 months ago

yes indeed , the correct order is :  
Step 1: #bin/bash -  
Step 2 -az webapp create  
Step 3- az webapp config hostname  
Step 4 -az webapp confing container set  
upvoted 57 times

🗨️ 👤 **Juanlu** 4 years, 6 months ago

Totally agree !  
upvoted 1 times

🗨️ 👤 **TechieBloke** 5 years, 4 months ago

Regimiento is right. Even the provided link show the same.  
Have to create the webapp before the config.  
upvoted 3 times

🗨️ 👤 **Regimiento** 5 years, 5 months ago

The answer has incorrect order:  
you first have to create web app, then you can config it.  
Therefore step 2 : az webapp create  
step 3: az webapp config hostname add  
upvoted 14 times

🗨️ 👤 **Nick\_1** 4 years, 11 months ago

The answer is correct as listed. The question does not mention that the image resides on local container, you need to upload to the ACR, hence you have that tag :v1  
In the same script the use the variable to define the webapp name, hence you can use this until the webapp is created, step 3  
upvoted 1 times

You develop a serverless application using several Azure Functions. These functions connect to data from within the code.  
You want to configure tracing for an Azure Function App project.  
You need to change configuration settings in the host.json file.  
Which tool should you use?

- A. Visual Studio
- B. Azure portal
- C. Azure PowerShell
- D. Azure Functions Core Tools (Azure CLI)

**Suggested Answer: B**

The function editor built into the Azure portal lets you update the function.json file and the code file for a function. The host.json file, which contains some runtime- specific configurations, is in the root folder of the function app.

```
FunctionApp
| - host.json
| - Myfirstfunction
| | - function.json
| | - ...
| - mysecondfunction
| | - function.json
| | - ...
| - SharedCode
| - bin
```

References:



<https://docs.microsoft.com/en-us/azure/azure-functions/functions-reference#fileupdate>

  **Chandu440**  4 years, 12 months ago




I think the answer is correct. The question talks about function app. A function app can only be created via portal or CLI. What we create in Visual studio is an Azure function that needs to be deployed within a function app. Function app is more like a container that can host many functions.  
upvoted 9 times

  **Juanlu** 4 years, 6 months ago

Agree !  
upvoted 1 times

  **quokka** 4 years, 9 months ago

That's misleading! Please try before you write. You can create function in both portal and in VS (and VSCode). If you change in portal, the code stored in storage account will not have version control (and lost if someone redeploys). If you want tracing only for temporal debugging, then you can do that in the portal. Otherwise best practice is to change in VS/VSCode with source control.  
upvoted 3 times

  **Kumar1988**  4 years, 2 months ago



Azure portal  
upvoted 1 times

  **Kumar1988** 4 years, 2 months ago

Azure portal  
upvoted 1 times

  **samyoyoyo** 4 years, 5 months ago

This is possible on the Azure portal only if the function App is running on Windows. In this case, we could use kudu to edit the host.json file. Otherwise, I think the most likely answer is to use Azure Functions core tool on Visual Studio Code  
upvoted 1 times

  **MyLife** 4 years, 6 months ago

The Functions editor built into the Azure portal lets you update your code and your function.json file directly inline.  
This is recommended only for small changes or proofs of concept - best practice is to use a local development tool like VS Code.

I would say portal.. as just host file changes.. correct me if i am wrong.

upvoted 1 times

🗨️ 👤 **Larry88** 4 years, 8 months ago

considering the requirement "You want to configure tracing for an Azure Function App project." A VS might be better answer.

upvoted 1 times

🗨️ 👤 **Steve\_az** 4 years, 9 months ago

How to edit functions in the Azure portal :

The Functions editor built into the Azure portal lets you update your code and your function.json file directly inline. This is recommended only for small changes or proofs of concept - best practice is to use a local development tool like VS Code.

<https://docs.microsoft.com/en-us/azure/azure-functions/functions-reference>

upvoted 2 times

🗨️ 👤 **Iloves73g** 5 years ago

Portal.

Is visual studio a must for us to develop az functions? No.

upvoted 2 times

🗨️ 👤 **Mnair** 5 years ago

Hey which is correct here ? portal or Visual studio ? i thought Visual studio.. thoughts please?

upvoted 1 times

🗨️ 👤 **mp007** 4 years, 11 months ago

Visual studio

upvoted 1 times

🗨️ 👤 **riley5** 5 years, 1 month ago

I thought portal also and then read this in ms documentation: "The Functions editor built into the Azure portal lets you update your code and your function.json file directly inline. This is recommended only for small changes or proofs of concept - best practice is to use a local development tool like VS Code." <https://docs.microsoft.com/en-us/azure/azure-functions/functions-reference#fileupdate>. So which one is the correct answer?

upvoted 1 times

🗨️ 👤 **riley5** 5 years, 1 month ago

So both are correct, but VScode is the most correct answer?

upvoted 1 times

🗨️ 👤 **lau13** 5 years ago

so, answer is either Portal or Visual Studio Code, but the option in the question is Visual Studio, not Visual Studio Code, that's why Portal is the correct answer.

upvoted 3 times

🗨️ 👤 **JaQua** 5 years, 3 months ago

they clearly state function app 'PROJECT' so visual studio is correct answer

upvoted 3 times

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

It is still project even if it is portal. you don't need to enable tracing locally. Just enable it in portal. Do your tracing and disable it.

upvoted 4 times

🗨️ 👤 **TechieBloke** 5 years, 4 months ago

You can use both but Portal is more convenient.

The function editor built into the Azure portal lets you update the function.json file and the code file for a function. The host.json file, which contains some runtime-specific configurations, is in the root folder of the function app.

upvoted 4 times

🗨️ 👤 **TaiH** 5 years, 4 months ago

Hei,

We can change the config in both Visual Studio and Azure Portal. Why do we choose Azure Portal here?

Tai

upvoted 1 times

You are developing a mobile instant messaging app for a company.

The mobile app must meet the following requirements:

- ⇒ Support offline data sync.
- ⇒ Update the latest messages during normal sync cycles.

You need to implement Offline Data Sync.

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

NOTE: Each correct selection is worth one point.

- A. Retrieve records from Offline Data Sync on every call to the PullAsync method.
- B. Retrieve records from Offline Data Sync using an Incremental Sync.
- C. Push records to Offline Data Sync using an Incremental Sync.
- D. Return the updatedAt column from the Mobile Service Backend and implement sorting by using the column.
- E. Return the updatedAt column from the Mobile Service Backend and implement sorting by the message id.

**Suggested Answer: BE**

B: Incremental Sync: the first parameter to the pull operation is a query name that is used only on the client. If you use a non-null query name, the Azure Mobile

SDK performs an incremental sync. Each time a pull operation returns a set of results, the latest updatedAt timestamp from that result set is stored in the SDK local system tables. Subsequent pull operations retrieve only records after that timestamp.

E (not D): To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause.

References:

<https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-offline-data-sync>

  **Az203test**  5 years, 1 month ago



Correct answer: B and D

upvoted 24 times

  **Jay\_2pt0** 5 years, 1 month ago

I believe you are right. <https://docs.microsoft.com/en-us/azure/app-service-mobile/app-service-mobile-offline-data-sync>

upvoted 3 times

  **riley5** 5 years, 1 month ago




I think you are perhaps right as well.

upvoted 1 times

  **Mvii** 5 years ago

Sorting by messageid column wouldn't sort by latest - sort by updatedAt descending instead. D makes more sense than E.

upvoted 2 times

  **AK89**  5 years ago

Correct answer: B and D

Link : <https://docs.microsoft.com/en-us/previous-versions/azure/app-service-mobile/app-service-mobile-offline-data-sync>

Phrase for clarification : "To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. "

upvoted 6 times

  **Kumar1988**  4 years, 2 months ago

How its option D , it should be E because sort by latest message id and get in offline sync

upvoted 1 times

  **silentsouls** 4 years, 3 months ago

Message-Id is usually even more convoluted then a guid.

Sorting on that would be a bad thing. so E is false.

D is correct.

upvoted 1 times

  **Larry616** 4 years, 9 months ago

Correct answer is B and D, coz it clearly said sorting by "updatedAt" column and doesn't support its own orderBy clause as below -

To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause.

upvoted 1 times

  **bluego78** 4 years, 9 months ago

Answers are B and D... (not E) !! You can't order a query in an offline data async operation, because the query "is always" and "must be" ordered by "updatedAt" field.

upvoted 2 times

  **ChandraShekhar** 4 years, 11 months ago

Hi,

since the UpdateAt is already sorted we can't add our sorting.

"To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. However, since the SDK adds its own sort on the updatedAt field, you cannot use a pull query that has its own orderBy clause."

So Answer should be B and E

upvoted 3 times

  **Anjalianju** 4 years, 11 months ago

are we sure is it B and D?

upvoted 1 times

  **malawneh** 5 years, 1 month ago

the latest updatedAt timestamp from that result set is stored in the SDK local system tables. Subsequent pull operations retrieve only records after that timestamp.

updatedAt is a column inside the local table, so the correct answer is B&D

<https://docs.microsoft.com/en-us/previous-versions/azure/app-service-mobile/app-service-mobile-offline-data-sync>

upvoted 5 times

  **Jedik** 5 years ago



"To use incremental sync, your server must return meaningful updatedAt values and must also support sorting by this field. " field/column - i guess

upvoted 2 times

  **Nids087** 5 years, 1 month ago

Everywhere I see answer as B and E. Need more explanation as to by it is D and not E

upvoted 1 times

  **riley5** 5 years, 1 month ago

Can you please explain why it's not b and e?

upvoted 1 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****Answer Area**

Strong

Bounded Staleness

Return the most recent patient status.

Consistent Prefix

Eventual

Return health monitoring data that is no less than one version behind.

After patient is discharged and all changes are assessed, retrieve the correct billing data with the final charges

**Suggested Answer:****Consistency levels****Answer Area**

Strong

Bounded Staleness

Return the most recent patient status.

Strong

Consistent Prefix

Eventual

Return health monitoring data that is no less than one version behind.

Bounded Staleness

After patient is discharged and all changes 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.

**References:**

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels>

BoundedStaleness

Eventual

upvoted 42 times

  **sathish11**  4 years, 10 months ago

Strong

BoundedStaleness

Eventual

upvoted 10 times

  **ss22**  4 years, 3 months ago

Correct Answer is " Strong , BoundedStaleness, ConsistentPrefix"

upvoted 1 times

  **palashroy** 4 years, 7 months ago

1. Strong, 2. Bounded staleness

3. Consistent Prefix: Final charges are sum of all charges. Eventual does not guarantee consistent read, thus reads might miss latest writes.



Also, all the writes history would be required for billing, thus Consistent Prefix guarantees always reads never sees out of order writes.

upvoted 1 times

  **luppittegui** 4 years, 5 months ago

Maths:  $a+b = b+a$

upvoted 2 times

  **NKnab** 4 years, 10 months ago

Correct answer shown by exam topics

upvoted 2 times

  **Pavan2358** 4 years, 11 months ago

what is the correct answer please?



I think its strong, bounded stateless and Consistent Prefix

upvoted 2 times

  **jv10** 5 years ago

what is the correct answer?

upvoted 4 times

  **vova12345** 5 years ago

strong, prefix, eventual

upvoted 5 times

  **longcui** 5 years, 1 month ago

would think the third answer is "session"

upvoted 1 times

  **saeza100** 5 years, 1 month ago

I think "Eventual,Bounded,Strong"

upvoted 2 times

## HOTSPOT -

You are creating a CLI script that creates an Azure web app related services in Azure App Service. The web app uses the following variables:

Variable name	Value
\$gitrepo	<a href="https://github.com/Contos/webapp">https://github.com/Contos/webapp</a>
&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:

**Suggested Answer:** *Explanation*

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.

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

References:

<https://medium.com/@satish1v/devops-your-way-to-azure-web-apps-with-azure-cli-206ed4b3e9b1>

 **hertino**  4 years, 9 months ago

Today I got 900, I had this question, my answer:

az appservice plan create


az webapp create

--plan \$webappname

az webapp deployment -

--repo-url \$gitrepo --branch master --manual-integration

upvoted 19 times

 **Kainite** 4 years, 9 months ago

Answer here: <https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-github?toc=/cli/azure/toc.json>

upvoted 9 times

 **gssicftohbrcsgtjg**  5 years, 3 months ago

Options

<https://www.examttopics.com/assets/media/exam-media/02521/0009700001.png>

upvoted 17 times

 **Nick\_1**  4 years, 11 months ago

Box 2 should be webapp create. Also, webapp name and plan name are the same, confusing...

upvoted 2 times

 **nikos2001** 4 years, 11 months ago



The first option in all selection boxes should be "az webapp create", not just "azure webapp" ?

# Create a web app.

```
az webapp create --name $webappname --resource-group myResourceGroup --plan $webappname
```

upvoted 7 times



  **crazymagi** 5 years, 1 month ago

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-github?toc=/cli/azure/toc.json>

upvoted 1 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><div>HTTP request header</div><div>Client cookie</div><div>HTTP message body</div><div>URL query string</div></div>
Encoding type	<div><div></div><div>HTML</div><div>URL</div><div>Unicode</div><div>Base64</div></div>

**Answer Area**

Property Value

Suggested Answer:

Client certificate location	<div><div></div><div>HTTP request header</div><div>Client cookie</div><div>HTTP message body</div><div>URL query string</div></div>
Encoding type	<div><div></div><div>HTML</div><div>URL</div><div>Unicode</div><div>Base64</div></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.


References:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-configure-tls-mutual-auth>

 **Avadhut01** Highly Voted 4 years, 11 months ago



Its correct answer.

upvoted 11 times

 **AshwaniT** Most Recent 4 years, 2 months ago



Correct.

upvoted 1 times

  **Sreddi** 4 years, 5 months ago

Correct answer

upvoted 1 times

  **GCMan** 4 years, 6 months ago

Correct.

upvoted 2 times

DRAG DROP -

You are developing a .NET Core model-view controller (MVC) application hosted on Azure for a health care system that allows providers access to their information.

You develop the following code:

```
services.AddAuthorization(options=>
{
    options.AddPolicy ("ProviderPartner", policy =>
    {
        policy.AddAuthenticationSchemes("Cookies, Bearer");
        policy.RequireAuthenticatedUser();
        policy.RequireRole("ProviderAdmin", "SysAdmin");
        policy.RequireClaim("editor", "partner");
    });
});
```

You define a role named SysAdmin.

You need to ensure that the application meets the following authorization requirements:

- ⇒ Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.
- ⇒ Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

How should you complete the code? To answer, drag the appropriate code segments to the correct locations. 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:

#### Suggested Answer: Explanation

Box 1:

Allow the ProviderAdmin and SysAdmin roles access to the Partner controller regardless of whether the user holds an editor claim of partner.

Box 2:

Limit access to the Manage action of the controller to users with an editor claim of partner who are also members of the SysAdmin role.

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

correct answer is:

1. [Authorize(Role="SysAdmin, ProviderAdmin")]
2. [Authorize(Policy="ProviderPartner")] [Authorize(Role="SysAdmin")]

upvoted 42 times

**igorcode** 5 years, 3 months ago

There is no ProviderPartner.. your answer is wrong.

upvoted 1 times

**chaudh** 4 years, 10 months ago

No,

[Authorize(Role="SysAdmin, ProviderAdmin")] is not provided in question

[Authorize(Role="SysAdmin", "ProviderAdmin")] is provided in question.

upvoted 2 times

**bhushan\_786** 4 years, 5 months ago

It is a typo from him while writing the answer. Look like his answer is correct.

upvoted 1 times

**chintan4190** 5 years, 1 month ago

This is correct answer.

<https://docs.microsoft.com/en-us/aspnet/core/security/authorization/roles?view=aspnetcore-3.1>

If you apply multiple attributes then an accessing user must be a member of all the roles specified

upvoted 5 times

**simonverma** 4 years, 1 month ago

1 is not correct. 2 is correct.

Either you write as [Authorize(Roles="SysAdmin, ProviderAdmin")] -- Roles- plural

or [Authorize(Role="SysAdmin, ProviderAdmin")] [Authorize(Role="SysAdmin, ProviderAdmin")].

upvoted 1 times

🗨️ 👤 **lywwind** Highly Voted 5 years, 4 months ago

Box 1:

```
[Authorize(Role= "ProviderAdmin")]
```

```
[Authorize(Role = "SysAdmin")]
```

Box 2:

```
[Authorize(Policy = "ProviderEditor", Role = "SysAdmin")]
```

Check this

<https://www.examttopics.com/exams/microsoft/az-202/view/4/>

upvoted 34 times

🗨️ 👤 **lau13** 5 years ago

The suggested box 1 means that it will only authorize someone who is both ProviderAdmin and SysAdmin at the same time.

upvoted 7 times

🗨️ 👤 **sinh** Most Recent 4 years, 7 months ago

`[Authorize(Role= "ProviderAdmin")][Authorize(Role = "SysAdmin")]` are true if both are satisfied.

`[Authorize(Role = "SysAdmin, ProviderAdmin")]` is true if either is met.

`[Authorize(Role = "SysAdmin, ProviderAdmin")]` is the correct answer, as we only need one of these roles here.

upvoted 3 times

🗨️ 👤 **Cornholioz** 4 years, 5 months ago

Isn't it the other way around? Requirement says Allow ProviderAdmin AND SysAdmin roles to Partner Controller. For AND, the box 1 should have it in two lines. So should the box 2?

Unsure!

upvoted 2 times

🗨️ 👤 **RahulKate** 4 years, 4 months ago

Options not mentioned with proper syntax,

```
[Authorize(Roles = "Administrator, PowerUser")]
```

so answer can be `[Authorize(Role= "ProviderAdmin")][Authorize(Role = "SysAdmin")]`

upvoted 1 times

🗨️ 👤 **Steve\_az** 4 years, 9 months ago

Answer option - [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_058\\_Image\\_0002.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_058_Image_0002.jpg)

upvoted 11 times

🗨️ 👤 **rafapaz09** 4 years, 9 months ago

Correct answer should be

Box 1.

```
[Authorize(Role="SysAdmin, ProviderAdmin")]
```

2.

```
[Authorize(Policy="ProviderPartner")]
```

```
[Authorize(Role="SysAdmin")]
```

Reference: <https://docs.microsoft.com/en-us/aspnet/core/security/authorization/roles?view=aspnetcore-3.1>

"If you apply multiple attributes then an accessing user must be a member of all the roles specified; the following sample requires that a user must be a member of both the PowerUser and ControlPanelUser role." By multiple attributes it means 2 different authorize attributes.

upvoted 6 times

🗨️ 👤 **Suchismita** 4 years, 10 months ago

Correct answers are

Box 1. (Allow the ProviderAdmin and SysAdmin roles access to the Partner controller)

```
[Authorize(Role= "ProviderAdmin")]
```

```
[Authorize(Role = "SysAdmin")]
```

Box2. (users with an editor claim of partner who are also members of the SysAdmin role that means user should be both Provider Partner and SysAdmin)

```
[Authorize(Policy="ProviderPartner")]
```

```
[Authorize(Role="SysAdmin")]
```

Refer <https://docs.microsoft.com/en-us/aspnet/core/security/authorization/roles?view=aspnetcore-3.1>

upvoted 5 times

🗳️ 👤 **Ave** 4 years, 6 months ago

This is the correct answer

upvoted 1 times

🗳️ 👤 **Not\_Sure\_Who\_Am\_I** 4 years, 11 months ago

Options: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_058\\_Image\\_0002.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_058_Image_0002.jpg)

upvoted 11 times

🗳️ 👤 **AK89** 5 years ago

Box 1 :

```
[Authorize(Role= "ProviderAdmin")]
```

```
[Authorize(Role = "SysAdmin")]
```

Box 2 :

```
[Authorize(Policy = "ProviderEditor", Role = "SysAdmin")]
```

upvoted 9 times

🗳️ 👤 **chintan4190** 5 years, 1 month ago

Options: <https://www.certlibrary.com/exam/AZ-203>

upvoted 15 times

🗳️ 👤 **Daltonic75** 5 years, 2 months ago

<https://docs.microsoft.com/en-us/aspnet/core/security/authorization/roles?view=aspnetcore-3.1>

"If you apply multiple attributes then an accessing user must be a member of all the roles specified; the following sample requires that a user must be a member of both roles."

The question said "Allow the ProviderAdmin and SysAdmin roles access". Both roles, so I think that is:

```
[Authorize(Role= "ProviderAdmin")]
```

```
[Authorize(Role = "SysAdmin")]
```

upvoted 4 times

🗳️ 👤 **hmirko** 5 years, 1 month ago

I think it states that users with any of those roles are allowed, correct answers are

1.

```
[Authorize(Role="SysAdmin, ProviderAdmin")]
```

2.

```
[Authorize(Policy="ProviderPartner")]
```

```
[Authorize(Role="SysAdmin")]
```

upvoted 20 times

🗳️ 👤 **oxaytol** 4 years, 11 months ago

this is the correct answer.

1. Must either have sysadmin or provider admin, so `[Authorize(Role="SysAdmin, ProviderAdmin")]`.

2. Must have and editor claim of partner AND ALSO be a member of sysadmin, so in order to require BOTH:

```
[Authorize(Policy="ProviderPartner")]
```

```
[Authorize(Role="SysAdmin")]
```

upvoted 6 times

🗳️ 👤 **shanky\_007** 5 years, 5 months ago

Box1. `[Authorize(Role = "SysAdmin, ProviderAdmin")]`

Box2. `[Authorize(Policy = "ProviderEditor", Role = "SysAdmin")]`

upvoted 15 times

🗳️ 👤 **gido** 5 years, 4 months ago

Correct answer example found here:

<https://stackoverflow.com/questions/700166/allow-multiple-roles-to-access-controller-action>

upvoted 4 times

🗳️ 👤 **LTiwana** 4 years, 10 months ago

Seems correct. The difference between the below two statements is mention in MS docs and as per the question "Allow the ProviderAdmin and SysAdmin roles access" means users can be part of either group.



-----

```
[Authorize(Role = "SysAdmin, ProviderAdmin")]
```

```
[Authorize(Role= "ProviderAdmin")]
```

```
[Authorize(Role = "SysAdmin")]
```

upvoted 1 times

  **Ieshuva** 4 years, 5 months ago

correct.bacause this Role = "SysAdmin, ProviderAdmin")]] means either of these roles

and this Authorize(Policy = "ProviderEditor", Role = "SysAdmin") means need to have both policy and the role

upvoted 1 times

  **Regimiento** 5 years, 5 months ago

If someone wonders what should be in those boxes:

Box 1:

```
[Authorize(Policy = "ProviderEditor")]
```

```
[Authorize(Role = "SysAdmin")]
```

Box 2:

```
[Authorize(Policy = "ProviderEditor", Role = "SysAdmin")]
```

upvoted 14 times

DRAG DROP -

You manage several existing Logic Apps.

You need to change definitions, add new logic, and optimize these apps on a regular basis.

What should you use? To answer, drag the appropriate tools to the correct functionalities. 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:

### Answer Area

#### Tools

#### Functionality

#### Tool

Logic Apps Designer

Edit B2B workflows

Code View Editor

Edit definitions in JSON

Enterprise Integration Pack

Visually add functionality

Suggested Answer:

### Answer Area

#### Tools

#### Functionality

#### Tool

Logic Apps Designer

Edit B2B workflows

Enterprise Integration Pack

Code View Editor

Edit definitions in JSON

Code View Editor

Enterprise Integration Pack

Visually add functionality

Logic Apps Designer

Box 1: Enterprise Integration Pack

After you create an integration account that has partners and agreements, you are ready to create a business to business (B2B) workflow for your logic app with the Enterprise Integration Pack.

Box 2: Code View Editor -

To work with logic app definitions in JSON, open the Code View editor when working in the Azure portal or in Visual Studio, or copy the definition into any editor that you want.

Box 3: Logical Apps Designer -

You can build your logic apps visually with the Logic Apps Designer, which is available in the Azure portal through your browser and in Visual Studio.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-enterprise-integration-b2b> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-author-definitions> <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-overview>

 **hertino** Highly Voted 4 years, 9 months ago


Today I got 900, I had this question, my answer:

Edit B2B Enterprise integration

Edit JSON Code View Editor

visually add logic Apps Designer

upvoted 11 times

 **nischal123** Highly Voted 4 years, 10 months ago





its the correct answer

upvoted 5 times

  **Sreddi**  4 years, 5 months ago

correct answer

upvoted 1 times

  **GCMan** 4 years, 7 months ago

correct

upvoted 1 times

You are implementing an Azure API app that uses built-in authentication and authorization functionality.

All app actions must be associated with information about the current user.

You need to retrieve the information about the current user.

What are two ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. HTTP headers
- B. environment variables
- C. /.auth/me HTTP endpoint
- D. /.auth/login endpoint

**Suggested Answer: AC**

A: After App Service Authentication has been configured, users trying to access your API are prompted to sign in with their organizational account that belongs to the same Azure AD as the Azure AD application used to secure the API. After signing in, you are able to access the information about the current user through the

HttpContext.Current.User property.

C: While the server code has access to request headers, client code can access GET /.auth/me to get the same access tokens (

References:

<https://docs.microsoft.com/en-us/azure/app-service/app-service-web-tutorial-auth-aad> <https://docs.microsoft.com/en-us/sharepoint/dev/spfx/web-parts/guidance/connect-to-api-secured-with-aad>

 **Larry616**  4 years, 9 months ago

Answer is correct

upvoted 5 times

 **VishwanathPathak**  4 years, 3 months ago

Ans is correct !

upvoted 2 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:

**Suggested Answer:** *Explanation*

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

 **realbart**  5 years, 1 month ago

The scale rule looks different now and has an explicit checkbox to divide the counts by the number of instances. I assume that this also is the case for the old situation, so the threshold is 1000 \*per instance\*

Even though you shouldn't configure it like this in real life (2100 messages in total would lead to switching between 2 and 3 instances every five minutes) I believe the requested answers are:

## METRIC SOURCE

- storage queue
- [service bus queue]
- current resource
- storage queue (classic)

## CRITERIA

- message count
- [active message count]

## TIME GRAIN STATISTICS

- total
- maximum
- [average]
- count

## OPERATOR

- greater than
- greater than or equal to
- less than
- [less than or equal to]

## ACTION

- increase count by
- increase count to
- [decrease count by]
- decrease count to

upvoted 57 times

🗨️ **AK89** 5 years ago

This one is the correct answer

upvoted 6 times

🗨️ **zainhuss** 4 years, 11 months ago

in the question they say is greater than 1000, so why we don't choose less than?

upvoted 1 times

🗨️ **oxaytol** 4 years, 11 months ago

because greater than 1000 does not include equal to 1000, so in order to match the reverse configuration you must choose less than or equal to (1000)

upvoted 5 times

🗨️ **Juanlu** 4 years, 2 months ago

Agree. Same as <https://www.examttopics.com/exams/microsoft/az-300/view/17/>

upvoted 1 times

🗨️ **Daltonic75** Highly Voted 🏆 5 years, 2 months ago

Options: [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 28 times

🗨️ **meoukg** Most Recent 🕒 3 years, 2 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 1 times

🗨️ **Steve\_az** 4 years, 9 months ago

<https://vceguide.com/how-should-you-configure-the-scale-rule/>

upvoted 2 times

🗨️ **hertino** 4 years, 9 months ago

Today I got 900, I had this question, my answer:

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 16 times

🗨️ **Cornholioz** 4 years, 5 months ago

I'm wondering about ActiveMessage Count. The scenario says "unprocessed and valid queue messages". Does Unprocessed also fall under ActiveMessages? Maybe I'm overthinking it. I'll go with ActiveMessage Count :)

upvoted 1 times

🗨️ **Kuna\_Lambo** 4 years, 2 months ago

<https://docs.microsoft.com/pl-pl/azure/service-bus-messaging/service-bus-metrics-azure-monitor>

upvoted 1 times

🗨️ **saeza100** 5 years, 1 month ago

Box 3 should be average or count ?

upvoted 2 times

🗨️ **riley5** 5 years, 1 month ago

Box 3 : Average

Question specifies "average queue length"

upvoted 4 times

🗨️ **Cornholioz** 4 years, 5 months ago

Average is correct. Although many sites are suggesting count, the question asks the following: "when the average queue length (of unprocessed and valid queue

messages) is greater than 1000"

upvoted 2 times

🗨️ 👤 **riley5** 5 years, 1 month ago

Some other places where I've seen this question, have the following for boxes 4-5:

Box4: greater than or equal to

Box5: increase count by

upvoted 3 times

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

its mentioned "a new rule that will [continuously] scale down the App Service" to it should be Decrease count by.

upvoted 4 times

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

someone please post the options.

upvoted 1 times

🗨️ 👤 **sho1** 5 years, 3 months ago

Box 3 : Average

Question specifies "average queue length"

upvoted 20 times

🗨️ 👤 **VK\_Gladiator** 4 years, 4 months ago

This rule is for scaling up. I believe it doesn't have anything to do with scaling down.

upvoted 1 times

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

we need to scale down, so decrease the count of vm's

upvoted 1 times

🗨️ 👤 **AnujD** 5 years, 3 months ago

Why decrease count by?

upvoted 1 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:

**Suggested Answer:** *Explanation*

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
```

References:

<https://docs.microsoft.com/en-us/azure/app-service/scripts/cli-deploy-staging-environment>

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


[https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_065\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_065_Image_0001.jpg)  
upvoted 36 times

 **malawneh** Highly Voted 4 years, 9 months ago

I passed today with score 860, i got this in my exam.  
My answer was:

```
group create
service plan create
webapp create
webapp deployment slot
webapp deployment source
```

upvoted 25 times

 **GokuUI** 4 years, 7 months ago

on point!  
upvoted 1 times

 **Steve\_az** Most Recent 4 years, 9 months ago

<https://vceguide.com/how-should-you-complete-the-commands-3/>  
upvoted 3 times

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

Someone please post the options

upvoted 2 times

🗨️ 👤 **JulienYork** 4 years, 4 months ago

can't you google by yourself but asking from someone else?

You may not find someone to google for you all the time, copy some part of the question and google it.. This is all other people do when you ask for the options..

you do not need someone to search for you. You can do by yourself

upvoted 2 times

## DRAG DROP -

You have a web app named MainApp. You are developing a triggered App Service background task by using the WebJobs SDK. This task automatically invokes a function code whenever any new data is received in a queue.

You need to configure the services.

Which service should you use for each scenario? To answer, drag the appropriate services to the correct scenarios. Each service 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:

**Suggested Answer:** *Explanation*

Box 1: WebJobs -

A WebJob is a simple way to set up a background job, which can process continuously or on a schedule. WebJobs differ from a cloud service as it gives you get less fine-grained control over your processing environment, making it a more true PaaS service.

Box 2: Flow -

Incorrect Answers:



Azure Logic Apps is a cloud service that helps you schedule, automate, and orchestrate tasks, business processes, and workflows when you need to integrate apps, data, systems, and services across enterprises or organizations. Logic Apps simplifies how you design and build scalable solutions for app integration, data integration, system integration, enterprise application integration (EAI), and business-to-business (B2B) communication, whether in the cloud, on premises, or both.

References:

<https://code.msdn.microsoft.com/Processing-Service-Bus-84db27b4>

  **cenaz**  5 years, 2 months ago

[https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_067\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_067_Image_0001.jpg) for options  
upvoted 19 times

  **levo017** 5 years, 2 months ago

According to options, Flow is correct answer. Flow is within a 'DevOps environment'. Logic App is overkill  
upvoted 4 times

  **DanielNadjbabi** 4 years, 5 months ago

[https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_067\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_067_Image_0001.jpg)  
upvoted 2 times

  **gssicftohbrcsgtjg**  5 years, 3 months ago

Box1: WebJobs

Box2: LogicApps

upvoted 9 times

  **ss22**  4 years, 3 months ago

Box 1: Logic Apps, Box2: WebJobs

upvoted 4 times

  **Cornholioz** 4 years, 5 months ago

This question is just unfair and plain stupid. The documentation on the comparison has DevOps preferred for both Logic Apps and WebJobs (search the page for DevOps and you'll see why). So while processing queue data can be done by both Logic Apps and WebJobs, so can the management of code in DevOps. So it's a nonsensical question.

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

upvoted 1 times

  **YogeshSoni** 4 years, 7 months ago

Both are WebJobs.

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

upvoted 5 times

  **Gummad** 4 years, 7 months ago



both are webJobs , see the below url

<https://www.examttopics.com/discussions/microsoft/view/7900-exam-az-300-topic-3-question-5-discussion/>

upvoted 3 times

🗨️ 👤 **pl9977** 4 years, 7 months ago

Confused.Whats the right answer here?

upvoted 2 times

🗨️ 👤 **SilNilanjan** 4 years, 9 months ago

Complete question: <https://vceguide.com/which-service-should-you-use-for-each-scenario/>

upvoted 2 times

🗨️ 👤 **Suchismita** 4 years, 10 months ago

Given answer is correct.

upvoted 2 times

🗨️ 👤 **ChandraShekhar** 4 years, 11 months ago

Box 2: WebJobs

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

upvoted 8 times

🗨️ 👤 **chaudh** 4 years, 10 months ago

Undeniable, 2nd box is still WebJobs

upvoted 7 times

🗨️ 👤 **Juanlu** 4 years, 3 months ago

Box 1: WebJobs

Box 2: WebJobs

Same question for AZ-300:

<https://www.examttopics.com/discussions/microsoft/view/7900-exam-az-300-topic-3-question-5-discussion/>

upvoted 2 times

🗨️ 👤 **ChandraShekhar** 4 years, 11 months ago

Box 2: WebApps

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

upvoted 6 times

🗨️ 👤 **LTiwana** 4 years, 10 months ago

Clearly stated in the documentation.

upvoted 3 times

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

\*WebJobs

upvoted 1 times

🗨️ 👤 **AK89** 5 years ago

Given answer is correct

upvoted 1 times

🗨️ 👤 **alirasouli** 5 years ago

Flow is technically a SaaS solution and it is rarely an answer in this exam. Flow is self-service while Logic Apps is for advanced integration. In the new software architecture models, it is recommended to use WebJobs only when the same Azure DevOps environment is required as Web Apps. For queue data processing, Logic Apps is more likely a solution than WebJobs.

My answers:

[Logic Apps]

[WebJobs]

upvoted 4 times

🗨️ 👤 **Daltonic75** 5 years, 2 months ago

<https://www.examttopics.com/discussions/microsoft/view/7900-exam-az-300-topic-3-question-5-discussion/>

upvoted 7 times

🗨️ 👤 **riley5** 5 years, 1 month ago



Yes, both are webjobs. Go the link and see HS007's comment.

upvoted 6 times

  **mahangu** 5 years, 4 months ago

WebJobs

upvoted 9 times

  **EYIT** 5 years, 5 months ago

Box 2: Logic App

upvoted 3 times

## HOTSPOT -

A company is developing a mobile app for field service employees using Azure App Service Mobile Apps as the backend.

The company's network connectivity varies throughout the day. The solution must support offline use and synchronize changes in the background when the app is online app.

You need to implement the solution.

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:

**Suggested Answer:** *Explanation*

Box 1: `var todoTable = client.GetSyncTable<TodoItem>()`

To setup offline access, when connecting to your mobile service, use the method `GetSyncTable` instead of `GetTable` (example):

`IMobileServiceSyncTable todoTable = App.MobileService.GetSyncTable(); /`

Box 2: `await todoTable.PullAsync("allTodoItems",todo.Table.CreateQuery());`

Your app should now use `IMobileServiceSyncTable` (instead of `IMobileServiceTable`) for CRUD operations. This will save changes to the local database and also keep a log of the changes. When the app is ready to synchronize its changes with the Mobile Service, use the methods

`PushAsync` and `PullAsync` (example): `await App.MobileService.SyncContext.PushAsync(); await todoTable.PullAsync();`

References:

<https://azure.microsoft.com/es-es/blog/offline-sync-for-mobile-services/>

 **Daltonic75** Highly Voted 5 years, 2 months ago

Options: [https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_069\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_069_Image_0001.jpg)

upvoted 8 times

 **GCMAN** Highly Voted 4 years, 7 months ago

answer is correct

upvoted 8 times

 **NPoul\_92** Most Recent 4 years, 8 months ago


<https://vceguide.com/how-should-you-complete-the-code-segment-26/>

upvoted 3 times

 **MuhammedHaneesh** 5 years, 2 months ago

Post the options please

upvoted 2 times

 **ebadali** 5 years, 2 months ago

[https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001\\_Page\\_069\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2019/10/Microsoft-AZ-203-date-01-06-2019-00001_Page_069_Image_0001.jpg)

upvoted 19 times

A company is developing a solution that allows smart refrigerators to send temperature information to a central location.

The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location.

You need to complete the configuration.

Which Azure CLI or PowerShell command should you run?

**New-AzureRmServiceBusNamespace**  
 -ResourceGroupName fridge-rg  
 -NamespaceName fridge-ns  
 -Location fridge-loc

A.

**az servicebus namespace create**  
 - -resource-group fridge-rg  
 - -name fridge-ns  
 - -location fridge-loc

B.

**New-AzureRmResourceGroup**  
 -Name fridge-rg  
 -Location fridge-loc

C.

**New-AzureRmServiceBusQueue**  
 -ResourceGroupName fridge-rg  
 -NamespaceName fridge-ns  
 -Name fridge-q  
 -EnablePartitioning \$False

D.

**Suggested Answer: B**

# Create a Service Bus messaging namespace with a unique name. Example: namespaceName=myNameSpace\$RANDOM az servicebus namespace create --resource-group \$resourceGroupName --name \$namespaceName --location eastus

References:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli>

  **motu**  5 years, 3 months ago

It should be D, a namespace has already been created.

upvoted 37 times

  **codebrew** 5 years, 3 months ago

D is not even an answer though, its blank?

upvoted 3 times

  **Dann1112** 5 years, 3 months ago

Letters are below their respective answer

upvoted 3 times

  **bogdan\_alex** 5 years, 2 months ago

the answer's text is above the letter ;)

upvoted 1 times

  **Larry616** 4 years, 9 months ago

Where did it mention "namespace has been created" in the question?

upvoted 1 times

  **Juanlu** 4 years, 6 months ago

Agree, D is the correct one !

upvoted 1 times

  **Illoves73g**  5 years ago

"You have an existing Service Bus.", means that you namespace has already been created, you just need to add a queue .

So, D is correct!

upvoted 12 times

🗨️ 👤 **ozigee** 5 years ago

"You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location." It's not said that it has already been created. So, B  
upvoted 2 times

🗨️ 👤 **agoncal** 4 years, 11 months ago

Some dumps have that script: "You have an existing Service Bus."  
If that script exists on the question, then D, otherwise B  
upvoted 4 times

🗨️ 👤 **Juanlu** 4 years, 2 months ago

Exactly, I agree. Nice comment, thanks to clarify !  
upvoted 2 times

🗨️ 👤 **MMJ** Most Recent 4 years, 1 month ago

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-quickstart-cli>  
upvoted 1 times

🗨️ 👤 **Kumar1988** 4 years, 2 months ago

i think B option  
upvoted 1 times

🗨️ 👤 **and\_p01** 4 years, 5 months ago

Answer is D. Please read the question carefully:  
The solution must receive and store messages until they can be processed. You create an Azure Service Bus instance by providing a name, pricing tier, subscription, resource group, and location.  
[You need to complete the configuration]?  
Its mean namespace is already created and the next logical step is to Create Service Bus Queue.  
upvoted 3 times

🗨️ 👤 **NajamKhan** 4 years, 6 months ago

"D" is the answer.  
Explanation:  
Guys, where did they say namespace create? Read the question carefully and properly "You create an Azure Service Bus instance" That's clearly mentioned, so the answer is "D".  
upvoted 3 times

🗨️ 👤 **codeoptimus** 4 years, 6 months ago

The answer should actually be C:  
1. The resource group has already been created  
2. The servicebus namespace has been created in the resource group  
3. The next and logical thing is to create the queue with PS command "New-AzureRmServiceBusQueue"  
upvoted 1 times

🗨️ 👤 **mxml** 4 years, 6 months ago

You mean D. The letter is below the image it points to.  
upvoted 2 times

🗨️ 👤 **a\_Ri** 4 years, 8 months ago

How could B be correct?  
There is no such Azure location. (should be something like eastus, eastus2 etc)  
upvoted 1 times

🗨️ 👤 **rod6214** 4 years, 8 months ago

It shouldn't be D because question says into the requirements "pricing tier" selection and only you can configure it by creating a namespace  
upvoted 1 times

🗨️ 👤 **Abbieboy** 4 years, 9 months ago

options: <https://vceguide.com/which-azure-cli-or-powershell-command-should-you-run/>  
upvoted 2 times

🗨️ 👤 **baltik5** 4 years, 9 months ago

They are asking about azure CLI command. Only cli command in there is B. Other commands are Azure PowerShell.  
upvoted 1 times

🗨️ 👤 **SunnyAU** 4 years, 9 months ago

The question asks "Which Azure CLI or PowerShell command should you run?". So both CLI or Powershell answers are sufficient.

upvoted 2 times

  **navaneeth** 4 years, 9 months ago

where did they say namespace create. Read the question properly "You create an Azure Service Bus instance"

upvoted 1 times



  **bluego78** 4 years, 9 months ago

In this question is not clear if you "have just created" or "have to create" the service bus namespace....

If the service bus namespace is just created, the correct answer is CREATE A QUEUE.



If the service bus namespace isn't created yet, you have to create it before to create the queue.

upvoted 1 times

  **dkaps06** 4 years, 9 months ago

I would take it that the namespace is created. This is because the question asks to "complete the configuration", in which case D should be answer to create the queue

upvoted 1 times

  **mrjust** 4 years, 10 months ago

Read again and this time carefully. Where does it say a namespace has been created ? There are two versions of this question.

The given answer is correct, in another version of the question a namespace is already created.

again dont blindly questions to not be tweaked

upvoted 3 times

  **minukobba** 4 years, 10 months ago

A and B is same answer, isn't it?

upvoted 2 times

  **novice25** 4 years, 7 months ago

No A is wrong --namespacename is used while creating queue not service bus namespace


upvoted 1 times

  **nikos2001** 4 years, 10 months ago

You create an Azure Service Bus instance BY PROVIDING a name, pricing tier, subscription, resource group, and location

You don't provide pricing tier for queues, you do it for namespaces

upvoted 1 times

  **godsbane** 4 years, 10 months ago

When creating a new Service Bus Instance, you HAVE to create a namespace too. So the namespace is already created, which means the answer should be D. : New-AzureRmServiceBusQueue.

upvoted 2 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. validate-jwt
- B. jsonp
- C. authentication-certificate
- D. check-header

**Suggested Answer: A**

Add the validate-jwt policy to validate the OAuth token for every incoming request.

Incorrect Answers:

B: 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.

References:

<https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-protect-backend-with-aad>

🗨️ 👤 **meoukg** 3 years, 2 months ago

Got it on 03/2022, I chose A. validate-jwt

upvoted 1 times

🗨️ 👤 **SidSadhu** 4 years, 9 months ago

100% correct

upvoted 3 times

🗨️ 👤 **hertino** 4 years, 9 months ago

It's correct

upvoted 2 times

## DRAG DROP -

A company backs up all manufacturing data to Azure Blob Storage. Admins move blobs from hot storage to archive tier storage every month. You must automatically move blocks to Archive tier after they have not been accessed for 180 days. The path for any item that is not archived must be placed in an existing queue. This operation must be performed automatically once a month. You set the value of TierAgeInDays to 180. How should you configure the Logic App? To answer, drag the appropriate triggers or action blocks to the correct trigger or action slots. Each trigger or action block 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:

### Triggers and Action Blocks

**Insert Entity**

\*Table: processing

\*Entity: Path >

Show advanced options

**Tier blob**

\*Blob path: Path >

\*Blob Tier: Archive

**When there are messages in a queue**

\*Queue Name: processing

Connected to: tierStorageAccountConnection

**Recurrence**

\*Interval: 1

\*Frequency: Month

### Answer Area

↓

**Set tier age variable**

↓

**List blobs**

↓

**For each**

Scan all blobs in this folder

\* Select an output from previous steps: value >

✓ If true

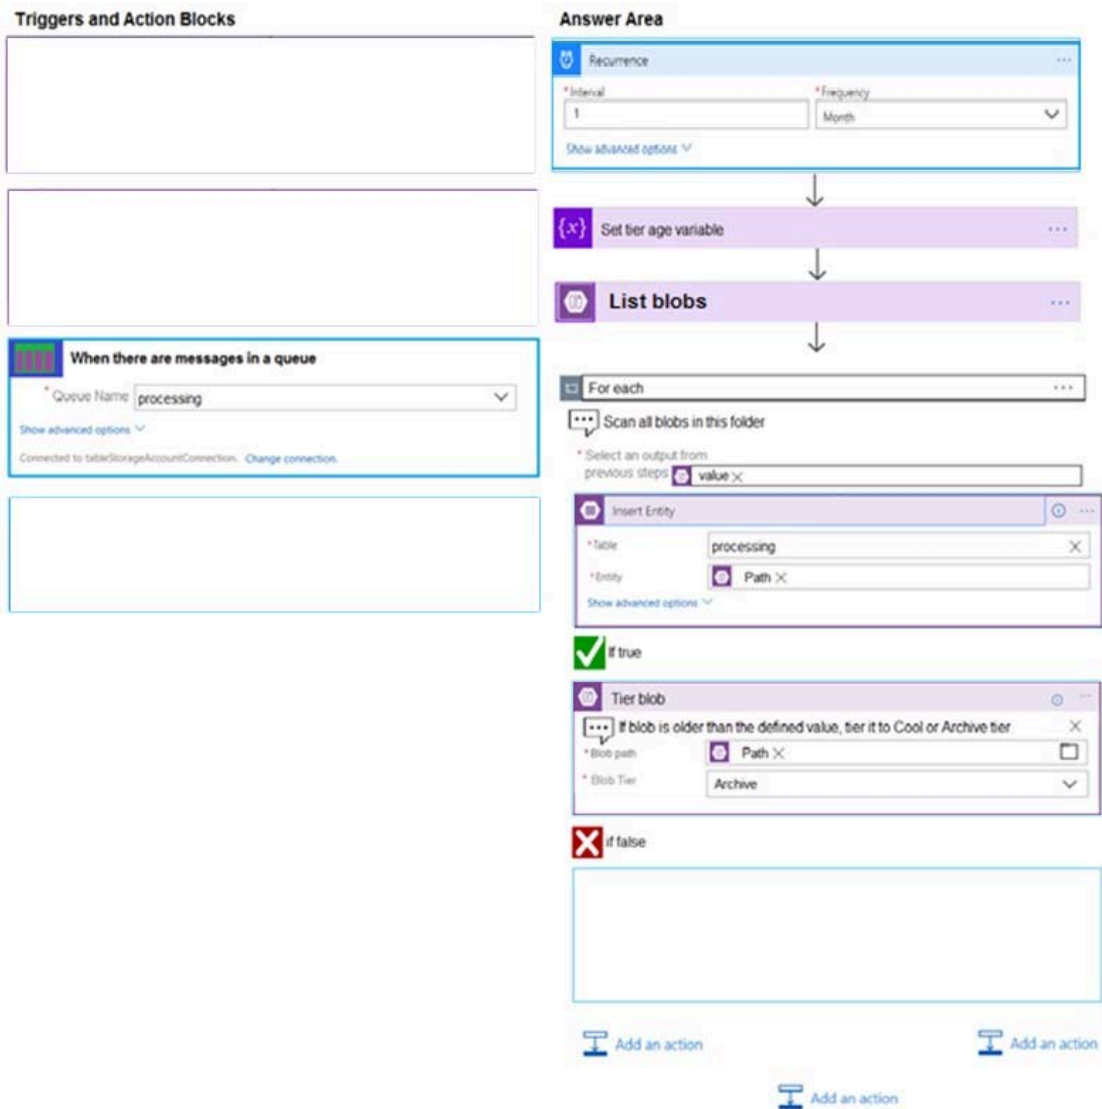
✗ If false

⚙ Add an action

⚙ Add an action



## Suggested Answer:



Box 1: Recurrence -

Box 2: Insert Entity -

Box 3 (if true): Tier Blob -

Box 4: (if false):

Leave blank.

References:

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations>

**Lkk51** Highly Voted 5 years ago

it should be

-> Recurrence

-> when there are msg in the queue

-> Tier blob

-> Insert entity

upvoted 28 times

**darijeb420** 4 years, 11 months ago

This would be correct because you would assume by using the queue it takes care of removing items that are processed of the queue, where as the answer given keeps inserting paths in the processing table

upvoted 1 times

**gunencali** 4 years, 5 months ago

I think your answer is wrong. Options are missing.

Explanation:

Box 1: Recurrence

This operation must be performed automatically once a month.

Box 2: Condition

Move blocks to Archive tier after they have not been accessed for 180 days.

Box 3 (if true): Tier Blob

Move blocks to Archive tier after they have not been accessed for 180 days.

Box 4: Put a message in a queue

The path for any item that is not archived must be placed in an existing queue.

References: <https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations>

upvoted 6 times

  **PriyaPT** 4 years ago

As per question, condition is already in place.

upvoted 1 times

  **Chilred**  4 years, 4 months ago

I think there are options missing ([https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020\\_Page\\_088\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020_Page_088_Image_0001.jpg)) or maybe this is different question

upvoted 10 times

  **Grancolrn** 4 years, 1 month ago

Definitely, have a look at Chilred's link :)

upvoted 2 times

  **DanielNadjbabi**  4 years, 5 months ago

Box 1: Recurrence

Box 2: Insert Entity

Box 3 (if true): Tier Blob

Box 4: (if false) Leave blank

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-perform-data-operations>

upvoted 3 times

  **SilNilanjan** 4 years, 8 months ago

How about -- 1) Recurrence, 2) Tier Blob, 3) Blank, 4) Insert entity

upvoted 1 times

  **azureuser1** 4 years, 10 months ago


1. Recurrence

2. Blank

3. Tier blob

4. Insert Entity - The path for any item that is not archived must be placed in an existing queue.

upvoted 4 times

  **chaudh** 4 years, 10 months ago

1. Recurrence - This operation must be performed automatically once a month

2. Tier blob - if blob is older than... - after they have not been accessed for 180 days.

3. Insert Entity - automatically move blocks to Archive tier

4. Leave blank (do nothing) - The path for any item that is not archived must be placed in an existing queue.

Note: message in a queue is not relevant, "processing" is Table, not Queue name

upvoted 1 times

  **nikos2001** 4 years, 11 months ago

What is the condition that refers to Yes / No boxes? None of the boxes in this question represent a condition ....

upvoted 2 times

  **oxaytol** 4 years, 11 months ago

This flow is poorly designed, how can this even be an answer?

upvoted 2 times

  **azurelearner666** 5 years ago

I believe this is incorrect.

the text says "The path for any item that is not archived must be placed in an existing queue." and the current solution is adding that for all the paths.  
so this action should be swapped with the Tier one that checks the 180 days condition and if so, places it on the Archive.

upvoted 1 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 an App Service that uses the Standard service tier. Configure the App service plan to automatically scale when the CPU load is high.
- D. Deploy the website to a virtual machine. Configure a Scale Set to increase the virtual machine instance count when the CPU load is high.

**Suggested Answer: C**

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.

chaudh Highly Voted 4 years, 10 months ago

given answer is correct, VM or shared service tier doesn't give availability.  
upvoted 7 times

Larry616 Highly Voted 4 years, 9 months ago

VM also works but it's more expensive than app service, the question requires minimizing cost.  
upvoted 5 times

gunencali Most Recent 4 years, 5 months ago

Is correct answer.

C . 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.  
upvoted 1 times

azurelearner666 5 years ago

Technically D would be also a valid solution, see <https://docs.microsoft.com/en-us/azure/virtual-machine-scale-sets/overview>  
On another side you would need a load balancer or an app gateway so it is not a complete response, but the statement is valid though.  
As usual, Exam questions are misleading ;).  
upvoted 1 times

Drunken 4 years, 11 months ago

Well the question is - what should you do ? Not which option would work.

and having VM with load balancer is just too much

upvoted 2 times

Umar93 4 years, 2 months ago

You want to minimize cost also.

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/weapp>.

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"
```

New-AzWebAppSlot	-Name myResourceGroup -Location \$location
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -Location \$location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -ResourceGroupName myResourceGroup -Slot review
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"
```

New-AzWebAppSlot	-Name myResourceGroup -Location \$location
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -Location \$location -ResourceGroupName myResourceGroup -Tier Standard
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -Location \$location -AppServicePlan \$webappname -ResourceGroupName myResourceGroup
New-AzWebApp	
New-AzAppServicePlan	
New-AzResourceGroup	

New-AzWebAppSlot	-Name \$webappname -ResourceGroupName myResourceGroup -Slot review
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
```

The New-AzResourceGroup cmdlet creates an Azure resource group.

The New-AzAppServicePlan cmdlet creates an Azure App Service plan in a given location

The New-AzWebApp cmdlet creates an Azure Web App in a given a resource group

The New-AzWebAppSlot cmdlet creates an Azure Web App slot.

References:

<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>

🗄️ 👤 **hertino** Highly Voted 4 years, 9 months ago

It's correct

upvoted 9 times

🗄️ 👤 **meoukg** Most Recent 3 years, 2 months ago

Got it on 03/2022, I chose as below:

New-AzResourceGroup

New-AzAppServicePlan

New-AzWebApp

New-AzWebAppSlot

upvoted 1 times

🗄️ 👤 **SidSadhu** 4 years, 9 months ago

It's correct, group->appserviceplan->webapp create->slot create

upvoted 2 times

🗄️ 👤 **malawneh** 4 years, 9 months ago

I passed today with score 860, i got this in my exam.

My answer was:

Same answer mentioned here

upvoted 2 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).

References:

<https://azure.microsoft.com/sv-se/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/>

Weird question. What VM have to do with this? On-premise is not IaaS, Azure Web App is PaaS not SaaS...



upvoted 1 times

  **sien** 4 years, 11 months ago

When you're not on an consumption based plan but an app service plan, you can configure how much instances (VM's) it uses under the via de App Service Plan - Scale out option.



So the answer is correct.

upvoted 3 times

  **pk90351** 4 years, 12 months ago

Pricing tier should be isolation or standard?

upvoted 1 times

  **whoisthis** 4 years, 11 months ago


It should be isolated as the azure resources must be inside VNET. Only Isolated app service plans be inside VNET. Rest all have links created to VNET

upvoted 4 times

  **azurelearner666** 5 years ago

If the ASE (App Service Environment) was listed, that could be a valid response too. <https://docs.microsoft.com/en-us/azure/app-service/environment/intro>

upvoted 1 times

  **godsbane** 4 years, 10 months ago

<https://azure.microsoft.com/en-us/blog/announcing-app-service-isolated-more-power-scale-and-ease-of-use/#:~:text=By%20selecting%20the%20Isolated%20pricing,your%20secure%2C%20dedicated%20virtual%20networks.>

upvoted 1 times

  **godsbane** 4 years, 10 months ago

You need to have an App Service Environment in order to run an Isolated App Service Plan.

upvoted 7 times



## HOTSPOT -

A company develops a series of mobile games. All games use a single leaderboard service.

You have the following requirements:

- ⇒ Code should 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 an Azure Cosmos database. 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 database. (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.Generate.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "Smith")
06         TableOperators.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal,
07             "ssmith@contoso.com")
08     ));
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:

**Suggested Answer:** *Explanation*

Box 1: Yes -

Code for CosmosDB, example:

```
// Parse the connection string and return a reference to the storage account.
CloudStorageAccount storageAccount = CloudStorageAccount.Parse(
    CloudConfigurationManager.GetSetting("StorageConnectionString"));
// Create the table client.
CloudTableClient tableClient = storageAccount.CreateCloudTableClient();
// Retrieve a reference to the table.
CloudTable table = tableClient.GetTableReference("people");
// Create the TableOperation object that inserts the customer entity.
TableOperation insertOperation = TableOperation.Insert(customer1);
```

Box 2: No -

A new record will always be added as TableOperation.Insert is used, instead of TableOperation.InsertOrReplace.

Box 3: No -

No partition key is used.

Box 4: Yes -  
References:  
<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

🗨️ **hertino** Highly Voted 4 years, 9 months ago

Today I got 900, I had this question, my answer:

Code for CosmosDB Yes

Save score func No

Data partitioned No

Code store parametres Yes

upvoted 24 times

🗨️ **Juanlu** 4 years, 5 months ago

Agree !

upvoted 4 times

🗨️ **heero** Highly Voted 5 years, 2 months ago

<https://www.sharetest.com/images/AZ-202.files/image137.jpg>

upvoted 20 times

🗨️ **Rabi3** Most Recent 1 year, 6 months ago

Can someone please explain something to me. The question are all related to to ComosDB however, cosmos db client is nowhere mentioned in the code. All i see are function of the storage account library.

upvoted 2 times

🗨️ **jay158** 4 years, 1 month ago

Answer is

Code for CosmosDb Yes

Save Score func NO ( because TableOertaion.Insert only inserts, does not update data)

Data partitioned Yes ( This is basic property of Table API, data is stored as PartitionKey+RowKey, and CosmosDB partitions data automatically)

Data stored as Gameld, PlayerID NO ( Gameld will be stored but its name will be partitionkey in DB, and PlayerId shall be RowKey)

So finally

Yes

NO

Yes

NO

upvoted 2 times

🗨️ **altafpatel1984** 3 years, 5 months ago

It is not asked that it will be stored as playerid, gameid. It simply asks will they be stored or not?

upvoted 1 times

🗨️ **NPaul\_92** 4 years, 8 months ago

<https://vceguide.com/when-users-reach-a-new-high-score-the-system-will-save-the-new-score-using-the-savescore-function-below/>

upvoted 3 times

🗨️ **StoneZuk** 4 years, 8 months ago

there is no partition mentioned for game. So answer is No

upvoted 1 times

🗨️ **malawneh** 4 years, 9 months ago

I passed today with score 860, i got this in my exam.

My answer was:

Code for CosmosDB Yes

Save score func No

Data partitioned No

Code store parametres Yes

upvoted 5 times

🗨️ **SkyDream** 4 years, 10 months ago

gameid is not partitionkey in this example, some code snippet has a line to set partitionkey. If that exists then Yes else No.

upvoted 3 times

🗨️ 👤 **chintan4190** 5 years, 1 month ago

gameid is partition key in <https://www.examttopics.com/discussions/microsoft/view/16364-exam-az-203-topic-24-question-20-discussion/>

so Box 3 should be yes

upvoted 3 times

🗨️ 👤 **mouni\_S6** 5 years ago

I am still now convinced that gameid is partition key, whats with <https://www.examttopics.com/discussions/microsoft/view/16364-exam-az-203-topic-24-question-20-discussion/> ,its the same question, what am I missing here?

upvoted 1 times

🗨️ 👤 **mouni\_S6** 5 years ago

I am still not convinced that gameid is partition key, whats with <https://www.examttopics.com/discussions/microsoft/view/16364-exam-az-203-topic-24-question-20-discussion/> ,its the same question, what am I missing here?

upvoted 1 times

🗨️ 👤 **hmirko** 5 years, 2 months ago

gameid is partitionkey, it looks like Box 3 answer is Yes

upvoted 4 times

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

how come you believe gameid is partitionkey?

upvoted 1 times

🗨️ 👤 **hmirko** 5 years, 2 months ago

<https://www.examttopics.com/discussions/microsoft/view/16364-exam-az-203-topic-24-question-20-discussion/>

upvoted 2 times

🗨️ 👤 **jw\_duke** 4 years, 10 months ago

I think this question is different than the link provided above.. based on the code shown. the partition key is not assigned during the SaveScore funtion.

upvoted 1 times

🗨️ 👤 **t8203** 4 years, 10 months ago

The partition key is set when it is calling `var scoreRecord = new PlayerScore(...)`

upvoted 2 times

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

Post the options please..

upvoted 2 times

## HOTSPOT -

You are working for a company that designs mobile applications. They maintain a server where player records are assigned to their different games. The tracking system is new and in development.

The application uses Entity Framework to connect to an Azure Database. The database holds a Player table and Game table.

When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record.

The application will call CreatePlayerWithGame with the correct gameId and the playerId to start the process. (Line numbers are included for reference only.)

```

01. namespace ContosoCradt
02. {
03.     public class PlayerDbContext : DbContext
04.     {
05.         public PlayerDbContext() : base ("name-dBConnString") { }
06.         public DbSet<Player> Players { get ; set ; }
07.         public DbSet<Game> Games { get ; set ; }
08.         protected override void OnModelCreating(DBModelBuilder modelBuilder)
09.         {
10.             modelBuilder.Entity<Player>().HasMany(x => x.Games).WithMany (x => x.Players);
11.         }
12.     }
13.     internal class dbConfiguration : DbMigrationConfiguration<PlayerDbContext>
14.     {
15.         public dbConfiguration() . {AutomaticMigrationsEnabled = true ; }
16.     }
17.     public class mp
18.     {
19.         public void CreatePlayerWithGame(int playerId, int gameId) => AddPlayer(playerId, GetGame(gameId));
20.         public Game GetGame(int gameId)
21.         {
22.             using (var db = new PlayerDbContext())
23.             {
24.                 return db.Games.FirstOrDefault(x => x.GameId == gameId);
25.             }
26.         }
27.         public Player AddPlayer (int playerId, Game game)
28.         {
29.             using (var db = new PlayerDbContext())
30.             {
31.                 var player = new Player
32.                 {
33.                     PlayerId = playerId,
34.                     Games = new List<Game> {game },
35.                 };
36.                 db.Players.Add(player);
37.                 db.SaveChanges();
38.                 return player;
39.             }
40.         }
41.     }
42.     public class Player
43.     {
44.         public int PlayerId { get ; set; }
45.         public string PlayerName { get ; set; }
46.         public virtual List<Game> Games { get ; set; }
47.     }
48.     public class Game
49.     {
50.         public int GameId { get ; set ; }
51.         public string Title { get ; set; }
52.         public string Platform { get ; set; }
53.         public virtual List<Player> Players { get ; set; }
54.     }

```

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:

**Suggested Answer:** *Explanation*

Box 1: Yes -

Box 2: No -

Box 3: Yes -

Box 4: No -

Many-to-many relationships without an entity class to represent the join table are not yet supported. However, you can represent a many-to-many relationship by including an entity class for the join table and mapping two separate one-to-many relationships.

protected override void OnModelCreating(ModelBuilder modelBuilder)

```
{
    modelBuilder.Entity<PostTag>()
        .HasKey(t => new { t.PostId, t.TagId });
    modelBuilder.Entity<PostTag>()
        .HasOne(pt => pt.Post)
        .WithMany(p => p.PostTags)
        .HasForeignKey(pt => pt.PostId);
    modelBuilder.Entity<PostTag>()
        .HasOne(pt => pt.Tag)
        .WithMany(t => t.PostTags)
        .HasForeignKey(pt => pt.TagId);
}
```

  **brajen**  4 years, 10 months ago

Discussion is getting confuse..

upvoted 27 times

  **gssicftohbrcrsgtjg**  5 years, 3 months ago

options

<https://www.examttopics.com/assets/media/exam-media/02521/0004800001.jpg>

upvoted 23 times

  **yllkabaci**  1 year, 1 month ago

I have writed the code and here is my answer:


1. No because we get an exception `SqlException: Cannot insert explicit value for identity column in table 'Player' when IDENTITY_INSERT is set to OFF`. Indeed you can't set the Id, it will be autoincremented (if you don't set `identity_insert` to off)
2. No because :
  - player will not be inserted
  - if we don't specify `PlayerId` we will get an error on Game like below: `SqlException: Cannot insert explicit value for identity column in table 'Games' when IDENTITY_INSERT is set to OFF`.
3. No because the code will not insert any row.
4. Yes, many-to-many relationship will be created.

upvoted 1 times

  **altafpatel1984** 3 years, 5 months ago

Entity class and `HasMany/WithMany` shows many to many is supported, but since code inserts player with only one gameid, it seems many to one only.

upvoted 1 times

  **altafpatel1984** 3 years, 6 months ago

Y

N

N (if Gamels is not type mistake)

Y (since Player class has `List<Game>` property and Game class has `List<Player>` property, and since it also mentioned as `HasMany-WithMany` so this is many to many relationship.)

Many to many since

upvoted 1 times

  **altafpatel1984** 3 years, 5 months ago

I guess Gamels is typo mistake other wise code will not compile and hence will not at all resulting into all answers as No !! So answer would be Y

N Y Y



upvoted 1 times

  **altafpatel1984** 3 years, 5 months ago

My previous answer slightly not correct.

Here is correct: Y N N Y

upvoted 1 times

  **dancsita** 4 years, 3 months ago

Y N Y N is the correct answer


1. Y - the player record will be inserted successfully by the AddPlayer
- 2 N - there's no code where a new gameid would be added
- 3 Y - when referring to the game id, a new object is created, it will empty
4. N - many-many relationship has a different syntax (HasMany, WithMany)

upvoted 2 times

  **altafpatel1984** 3 years, 5 months ago



HasMany/WithMany is already there in code.

upvoted 1 times

  **cbn** 4 years, 3 months ago

Wonder why a complete programming question coming under Azure!

upvoted 2 times

  **minsma** 4 years, 3 months ago


because it is a Developer certificate..

upvoted 2 times

  **tmfahim** 4 years, 4 months ago

If the text for this question is correct, then everything should be "No". There is no property "Gameld" in the Game class, rather it is "Gamels". So the code will definitely not run.


upvoted 1 times

  **Juanlu** 4 years, 3 months ago

Agree:

1. YES
2. NO
3. YES. Gameld is not a property of Games, so GetGame(id) doesn't work. But if it a typo the answer is NO.
4. NO.

upvoted 3 times

  **dylansun** 4 years, 4 months ago

I think the answer should be Yes, No, No, No

upvoted 1 times

  **NajamKhan** 4 years, 6 months ago

guys, All are YES. The correct answer, tested on EF console application.

upvoted 6 times

  **swiftfoxmark2** 4 years, 7 months ago

I copied this code to an Entity Framework console application and ran it. The answers I got were all "Yes" in this case.

If using Entity Framework Core, however, the last answer would be "No".

upvoted 5 times

  **Mike204** 4 years, 8 months ago


Answers do not matter if you cannot view the questions.

upvoted 1 times

  **lamamo8656** 4 years, 8 months ago

<https://www.examttopics.com/assets/media/exam-media/02521/0004800001.jpg>

upvoted 1 times

  **SagarShete** 4 years, 9 months ago

According to me Answer is

Yes - Code will successfully insert the player record

No - There is no bug in the code and no additional copy will be created.



No - No getting why will it be wrong game id when i am already passing the gameid in function.

AddPlayer(playerId, GetGame(gameId)) -- Which will get the exact associated game id.

No - Yes, Entity framework supports many to many but it is in this context. Is it required. According to me it's not required.

The problem statement states " When adding a player, the code should insert a new player record, and add a relationship between an existing game record and the new player record."

upvoted 1 times



  **quokka** 4 years, 9 months ago

It's Y N N N.

The last N for Many-to-Many relationship should be like the example in [https://docs.microsoft.com/en-us/ef/core/modeling/relationships?](https://docs.microsoft.com/en-us/ef/core/modeling/relationships?tabs=fluent-api%2Cfluent-api-simple-key%2Csimple-key#many-to-many)



tabs=fluent-api%2Cfluent-api-simple-key%2Csimple-key#many-to-many

upvoted 2 times

  **quokka** 4 years, 9 months ago

I meant Y N N Y. Found a M-M relationship example from another MS doc <https://docs.microsoft.com/en-us/aspnet/mvc/overview/getting-started/getting-started-with-ef-using-mvc/creating-a-more-complex-data-model-for-an-asp-net-mvc-application#add-code-to-database-context>

upvoted 7 times

  **Dishak** 4 years, 9 months ago

Answer is Yes No No Yes

upvoted 6 times

  **Mr2302682** 4 years, 9 months ago

The Answer is:



1-Yes - why shouldn't it be inserted?

2-No - You read an existing game and put it into the player instance. This is just a reference!

3-No- See the answer before. It is just a reference. The player played a game that already existed.

4-Yes - But here I'm not sure.

upvoted 8 times

  **Kainite** 4 years, 9 months ago

This is confusing and I don't see any concrete justification so this is what I think is the correct Answer:

Yes - this inserts a player record into players, not a new player because to get the player context this player need to exist

No - it's not inserting new games and the id it gets is from one already exist on the players context

Yes - return db.games.FirstOrDefault (x=> x.GameId==gameId) i do believe it will get any other game not the one in context because is get the first game ID of the player in a context that may not be the correct one.

Yes - The class Player has public virtual List<Games> and the class Game has public virtual List<Players> so this means that need to have a M-M relationship

upvoted 1 times

  **Mr2302682** 4 years, 9 months ago

why do you think you get the wrong gameid? The GameId is unique!

upvoted 3 times

You use Azure Table storage to store customer information for an application. The data contains customer details and is partitioned by last name. You need to create a query that returns all customers with the last name Smith. Which code segment should you use?

- A. `TableQuery.GenerateFilterCondition("PartitionKey", Equals, "Smith")`
- B. `TableQuery.GenerateFilterCondition("LastName", Equals, "Smith")`
- C. `TableQuery.GenerateFilterCondition("PartitionKey", QueryComparisons.Equal, "Smith")`
- D. `TableQuery.GenerateFilterCondition("LastName", QueryComparisons.Equal, "Smith")`

**Suggested Answer: C**

Retrieve all entities in a partition. The following code example specifies a filter for entities where 'Smith' is the partition key. This example prints the fields of each entity in the query results to the console.

Construct the query operation for all customer entities where `PartitionKey="Smith"`.

```
TableQuery<CustomerEntity> query = new TableQuery<CustomerEntity>().Where(TableQuery.GenerateFilterCondition("PartitionKey",  
QueryComparisons.Equal,  
"Smith"));
```

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

  **Larry616**  4 years, 9 months ago

Answer is correct

upvoted 6 times

  **ss22**  4 years, 3 months ago

Given answer is correct


upvoted 1 times

  **bugimachi** 4 years, 4 months ago

Better link:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-design-guide>

upvoted 2 times

  **SidSadhu** 4 years, 9 months ago

Correct

upvoted 2 times



## HOTSPOT -

You are developing an app that manages users for a video game. You plan to store the region, email address, and phone number for the player. Some players may not have a phone number. The player's region will be used to load-balance data.

Data for the app must be stored in Azure Table Storage.

You need to develop code to retrieve data for an individual player.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Suggested Answer:** *Explanation*

Box 1: region -

The player's region will be used to load-balance data.

Choosing the PartitionKey.

The core of any table's design is based on its scalability, the queries used to access it, and storage operation requirements. The PartitionKey values you choose will dictate how a table will be partitioned and the type of queries that can be used. Storage operations, in particular inserts, can also affect your choice of PartitionKey values.

Box 2: email -

Not phone number some players may not have a phone number.

Box 3: CloudTable -

Box 4 : TableOperation query =..

Box 5: TableResult -

References:

<https://docs.microsoft.com/en-us/rest/api/storageservices/designing-a-scalable-partitioning-strategy-for-azure-table-storage>

 **realbart**  5 years, 1 month ago

Options: <https://www.examttopics.com/assets/media/exam-media/02521/0004300001.png>

upvoted 12 times

 **mp007**  4 years, 11 months ago

Given answer is correct.

upvoted 11 times

 **dylansun**  4 years, 4 months ago

The given answer is correct.

You can check the MS docs example: <https://docs.microsoft.com/en-us/azure/cosmos-db/tutorial-develop-table-dotnet#get-an-entity-from-a-partition>

upvoted 3 times

 **samyoyoyo** 4 years, 5 months ago

The correct answer is definitely :

Region

Email

CloudTable

TableOperation

TableResult

upvoted 7 times

 **subrata83** 4 years, 8 months ago


Given answer is correct . Refer <https://docs.microsoft.com/en-us/azure/cosmos-db/tutorial-develop-table-dotnet>

upvoted 2 times

 **NPoul\_92** 4 years, 8 months ago

<https://vceguide.com/how-should-you-complete-the-code-84/>

upvoted 2 times

  **AK89** 5 years ago

The correct answer is :

Region

Email

CloudTable

TableOperation

TableQuery



upvoted 6 times

  **superman917** 4 years, 11 months ago

The last option might be TableResult.

<https://vkinfotek.com/azureqa/how-do-i-retrieve-a-single-entity-from-table-storage-using-tableoperation.html>

upvoted 19 times

  **BakePVL** 4 years, 11 months ago

+1 for TableResult

Another example <https://www.c-sharpcorner.com/article/retrieving-records-understanding-table-storage-part-three/>

upvoted 3 times

  **MuhamedHaneesh** 5 years, 2 months ago

Options please

upvoted 2 times

## HOTSPOT -

You are developing a data storage solution for a social networking app.

The solution requires a mobile app that stores user information using Azure Table Storage.

You need to develop code that can insert multiple sets of user information.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Suggested Answer:** *Explanation*

Box 1, Box 2: TableBatchOperation

Create the batch operation.

```
TableBatchOperation op = new TableBatchOperation();
```

Box 3: ExecuteBatch -

/ Execute the batch operation.

```
table.ExecuteBatch(op);
```

Note: You can insert a batch of entities into a table in one write operation. Some other notes on batch operations:

You can perform updates, deletes, and inserts in the same single batch operation.

A single batch operation can include up to 100 entities.

All entities in a single batch operation must have the same partition key.

While it is possible to perform a query as a batch operation, it must be the only operation in the batch.

References:


<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-dotnet>

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

options

<https://www.examttopics.com/assets/media/exam-media/02521/0003700001.png>

upvoted 13 times

 **gematsaljoa** Highly Voted 4 years, 5 months ago

1,2 : TableBatchOperation

3 : ExecuteBatch

upvoted 8 times

 **clarionprogrammer** Most Recent 4 years, 1 month ago

Answer is correct.

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.cosmos.table.cloudtable.executebatch?view=azure-dotnet>

upvoted 1 times

 **chicagomassageseeker** 4 years, 1 month ago

Answers:

// Define a batch operation.

1. TableBatchOperation op = new TableBatchOperation();

2. // Execute the batch of operations on the "clients" table.

```
Table.execute(op);
```

Links: See Insert of batch of entities section:

<https://docs.microsoft.com/en-us/azure/cosmos-db/table-storage-how-to-use-java>

upvoted 3 times

 **chicagomassageseeker** 4 years, 1 month ago

Answers:

1. TableBatchOperation op = new TableBatchOperation();

2. // Execute the batch of operations on the "clients" table.

```
Table.execute(op);
```

upvoted 3 times

🗨️ 👤 **Bluediamond** 4 years, 10 months ago

wouldn't last box be insert?

upvoted 1 times

🗨️ 👤 **svcoco** 4 years, 7 months ago

Yes but that part of the code is omitted, so after Insert next method is ExecuteBatch. Answer is correct

<https://www.c-sharpcorner.com/article/batch-table-operations-understanding-table-storage-part-five/>

upvoted 1 times

You must implement Application Insights instrumentation capabilities utilizing the Azure Mobile Apps SDK to provide meaningful analysis of user interactions with a mobile app.

You need to capture the data required to implement the Usage Analytics feature of Application Insights.

Which three data values should you capture? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Trace
- B. Session Id
- C. Exception
- D. User Id
- E. Events




**Suggested Answer: ADE**

Application Insights is a service for monitoring the performance and usage of your apps. This module allows you to send telemetry of various kinds (events, traces, etc.) to the Application Insights service where your data can be visualized in the Azure Portal.

Application Insights manages the ID of a session for you.



References:

<https://github.com/microsoft/ApplicationInsights-Android>

  **LEOPOLD**  5 years, 1 month ago


The Events tool is the only Usage tool that does not require a user ID or a session ID. (<https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-overview>) So, B D E

upvoted 21 times

  **hmirko** 5 years, 1 month ago


Agree, <https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-troubleshoot>

upvoted 1 times

  **dery** 5 years, 1 month ago

"To track users over time, Application Insights requires a way to identify them. The Events tool is the only Usage tool that does not require a user ID or a session ID." so B and D should be there, <https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-overview>



upvoted 1 times

  **godsbane** 4 years, 10 months ago

It doesn't require an User ID or Session ID for tracking Events. But if you want to track what users are doing in your app, read here:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-send-user-context>. You need User Id and Session ID. Answer is Session Id, User ID, Events.

upvoted 9 times

  **LTiwana** 4 years, 10 months ago




Yes, B and D are needed. That's what Leopold's link shows and is asked in the question as well "You need to capture the data required to implement the Usage Analytics feature of Application Insights." Even if you just google this string from the question, Leopold's provided link shows up on top of the search results.

upvoted 1 times

  **Juanlu** 4 years, 3 months ago




I Agree !!

upvoted 1 times

  **AK89**  5 years ago

Correct Answer is : Events, UserID and Session ID

upvoted 17 times

  **ahadjithoma**  4 years, 5 months ago

B,D,E

In order to track what a user does over time, Application Insights needs an ID for each user or session. Include the following IDs in every custom event or page view.

Users, Funnels, Retention, and Cohorts: Include user ID.

Sessions: Include session ID.

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-send-user-context#tracking-users>

upvoted 1 times

🗨️ 👤 **luppitguei** 4 years, 5 months ago

Session, User and event: <https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-segmentation>

upvoted 1 times

🗨️ 👤 **samyoyoyo** 4 years, 5 months ago

Absolutely wrong, what's the matter with you ? The right answer is : Trace, Exception, Events.

upvoted 4 times

🗨️ 👤 **Cornholioz** 4 years, 5 months ago

Being offensive doesn't make you right. Providing explanations might.

To know "user interactions", how will Trace info and Exception info help?

User ID's and the related Sessions, and Events will help.

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-segmentation>

Anyone who answers this question in the most incorrect way will still get 33% of the marks allotted. That's the least you can score if you simply attempt this question. You seemed to have managed the most incorrect answer.

upvoted 1 times

🗨️ 👤 **Cornholioz** 4 years, 6 months ago

If you want analysis for All Users, then you don't need User ID... right?

Assuming User Id is not required, for usage analytics, would you need to scope it for Sessions? Why not for all Sessions, just so you get analytics for 'Usage' and not scoped for sessions?

Does it then make the given answer correct?

upvoted 1 times

🗨️ 👤 **chaudh** 4 years, 10 months ago

BDE is my choice

upvoted 1 times

🗨️ 👤 **oxaytol** 4 years, 11 months ago

Application Insights manages the ID of a session for you, that excludes B.

You need the userid (D). So my bet is given answer is correct.

upvoted 2 times

🗨️ 👤 **agoncal** 4 years, 11 months ago

B,D,E

"In order to track what a user does over time, Application Insights needs an ID for each user or session. Include the following IDs in every custom event or page view.

Users, Funnels, Retention, and Cohorts: Include user ID.

Sessions: Include session ID."

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/usage-send-user-context>

upvoted 3 times

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

given is correct

upvoted 12 times

🗨️ 👤 **PraPat** 5 years, 4 months ago

Ans is - A,C and E <https://docs.microsoft.com/en-us/azure/azure-monitor/app/asp-net-exceptions>

upvoted 4 times

🗨️ 👤 **lywwind** 5 years, 4 months ago

Not for exception,

upvoted 2 times

🗨️ 👤 **samyoyoyo** 4 years, 5 months ago



I confirm

upvoted 1 times

  **auma5545** 4 years, 5 months ago

given answer also same in whizlabs Q23 in 3rd

upvoted 1 times

  **altafpatel1984** 3 years, 5 months ago

Question is for Usage Analytics and you given answer for exception diagnostic. Please first read question properly.

upvoted 1 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	The 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	The subscriprion 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 subscription. 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:

**Suggested Answer:** *Explanation*

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.

References:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

 **oxaytol**  4 years, 11 months ago

Based on options you have: SQLFilter, CorrelationFilter or No Filter, the solution is:

FutureOrders: SQLFilter -> You need to specify a condition for not retrieving results. You cannot achieve this with correlation filter or no filter.

HighPriorityOrders -> Correlation filter as you can handle this with AND operations and the fields.



InternationalOrders -> SQL Filter because you cannot set a "NOT" instruction using correlation.

HighQuantityOrders: As above, you need a SQL filter in order to include a sentence with "greater than".

All orders: No filter, so you will retrieve all the information.

upvoted 29 times

  **brajen** 4 years, 10 months ago

For Future Order False-filter is the right choice but it's not given in the question so we can use SQL filter for the same.

High Priority Orders are having multiple condition so Correlation Filter is correct

International-orders: just one condition so use SQL Filter

All Order : means no filter required

upvoted 2 times

  **Juanlu** 4 years, 3 months ago

Totally Agree !

upvoted 1 times

  **hertino**  4 years, 9 months ago

Today I got 900, I had this question, my answer:

FutureOrders: SQLFilter

HighPriorityOrders: CorrelationFilter

InternationalOrders: SQLFilter

HighQuantityOrders: SQLFilter

AllOrders: No Filter

upvoted 23 times

  **pranay** 4 years, 8 months ago

If you remove the existing filter(sqlFilter/TrueFilter), you will not receive any message. So no filter cannot be an answer for AllOrder.

upvoted 2 times

  **cbn**  4 years, 3 months ago

This is what I think: (No Filter / SQL / SQL / SQL / SQL)

FutureOrders: [No Filter] (When there are no filters, no message will be delivered)

High Priority Orders: [SQL] (High priority AND (Region <> US))


International Orders: [SQL] (Region <> US)

High Quantity Orders: [SQL] (Quantity > 100)

All Orders: [SQL] (We need to define an action and only SQL filters support that)

It is a bad practice to use correlation filters in unintended ways. I don't think this exam expects such answers, even though it is technically possible.

upvoted 4 times

  **monniq** 4 years, 3 months ago



I agree. Went through the documentation and tried it on the portal and came to the same conclusion.

upvoted 3 times

  **sleepyboy777** 4 years, 3 months ago

Boolean filter is derived from SQL filter, so since boolean filter is not an option, answer is SQL filter. True is 1=1 and false is 1=0. You technically can use correlation filter and set one of the properties to an impossible value but this is bad practice because it is confusing and what if the value became valid/possible

upvoted 1 times

  **samch** 4 years, 4 months ago

FutureOrders: Would use a Boolean filter as to me it sounds conditional, as BooleanFilter isn't an option would choose SQLFilter.



HighPriorityOrders: CorrelationFilter - since correlation id is used and you should choose correlation filters over SQL-like filters because they're much more efficient in processing and have less impact on throughput

InternationalOrders: SQLFilter as its conditional. i.e. Country NOT USA.

HighQuantityOrders: Could be SQLFilter or CorrelationFilter since it's a user defined property and correlation filters are more efficient to use. Although, couldn't see any logic on doing a greater than in a correlation filter so a SQL Filter is needed.

AllOrders: No Filter - As it affects every single order, so no filters needed.

upvoted 1 times

  **N8n0z** 4 years, 7 months ago

FutureOrders -> CorrelationFilter (We need to specify a property so it won't match)

HighPriorityOrders -> CorrelationFilter(The CorrelationId is a measure of priority)

InternationalOrders -> SQLFilter(Use the region)

HighQuantityOrders -> CorrelationFilter(It's a user defined property)

AllOrders -> No Filter ("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")

All here: <https://docs.microsoft.com/pt-pt/azure/service-bus-messaging/topic-filters>

upvoted 1 times

🗳️ 👤 **NPaul\_92** 4 years, 8 months ago

<https://vceguide.com/which-filter-types-should-you-implement/>

upvoted 3 times

🗳️ 👤 **Larry616** 4 years, 9 months ago

CorrelationFilter (more efficient)

SQLFilter (Note high priority orders also contains international order)

SQLFilter

SQLFilter

NoFilter (Default to TrueFilter)

upvoted 5 times

🗳️ 👤 **quokka** 4 years, 9 months ago

FutureOrders: CorrelationFilter (FalseFilter is not an option here, so set equal to "xxx no match" to ensure won't match)

HighPriorityOrders: SQLFilter (high is relative, also need logical AND ShipLocation NOT 'US')

InternationalOrders: SQLFilter (ShipLocation NOT 'US')

HighQuantityOrders: SQLFilter (numeric comparison to 100 units)

AllOrders: No Filter (default to let all orders into subscription)

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/topic-filters>

Correlation filters - needs exact equal match (string is case-sensitive), there's no NOT operator, only AND.

SQL filters - to test against: exists, is null, NOT/AND/OR, relational, numeric arithmetic, text pattern LIKE).

Boolean filters - true or false, defaults to true and don't filter messages

upvoted 3 times

🗳️ 👤 **realbart** 4 years, 9 months ago

I cannot edit my earlier answer, so here is

This question talks about a queue with several subscriptions. Surely they mean a topic? (<https://www.serverless360.com/blog/azure-service-bus-queues-vs-topics>)

Also, removing a default Rule/Filter would cause no message to be delivered to the subscriber. So in my view, the answers would be:

FutureOrders: [No Filter]

High Priority Orders: [CorrelationFilter] (Priority == true can be done using correlation)

International Orders: [SQLFilter] (Correlation does not have a "does not equal" option)

High Quantity Orders: [SQLFilter] (Greater Than can be done using SQL Filter)

All Orders [ ] (You get a true filter by default. A SQL Filter with 1==1 might also do the job. This one is ambiguous)

As mentioned before, catching all or nothing can be achieved more explicitly selecting a TrueFilter or FalseFilter.

upvoted 1 times

🗳️ 👤 **harry973** 4 years, 10 months ago

high priority orders include international orders too. International orders are ones which are from other than US. So should it not be sql filters not correlation filter?

upvoted 5 times

🗳️ 👤 **chaudh** 4 years, 10 months ago

given answer is correct

upvoted 2 times

🗨️ 👤 **pranay** 4 years, 10 months ago

Answer is NoFilter, CorrelationFilter, SQLFilter, SQLFilter, SQLFilter

1. By default we have sqlfilter (whose condition is 1=1 means true). If we delete default sqlfilter and add no other filter, it will not receive any messages so answer is NoFilter

2. If more than 1 condition is there, correlationFilter should be preferred.

3. CorrelationFilter cannot check not match condition. SqlFilter can

4. sql expression so SQLFilter

5. default SQLFilter will work as its value is 1=1 and allows all messages to be received.

upvoted 7 times

🗨️ 👤 **brajen** 4 years, 10 months ago

Given answer is Correct only..

upvoted 1 times

🗨️ 👤 **nikos2001** 4 years, 11 months ago

I would say FutureOrders -> CorrelationFilter. You can specify values for Label or CorrelationId properties which do not exist in your messages  
CorrelationFilter Filter = new CorrelationFilter() { CorrelationId = "value\_not\_exist" };

[https://github.com/Azure/azure-service-](https://github.com/Azure/azure-service-bus/tree/master/samples/DotNet/GettingStarted/Microsoft.Azure.ServiceBus/TopicSubscriptionWithRuleOperationsSample)

[bus/tree/master/samples/DotNet/GettingStarted/Microsoft.Azure.ServiceBus/TopicSubscriptionWithRuleOperationsSample](https://github.com/Azure/azure-service-bus/tree/master/samples/DotNet/GettingStarted/Microsoft.Azure.ServiceBus/TopicSubscriptionWithRuleOperationsSample)

upvoted 1 times

🗨️ 👤 **justicetree555** 4 years, 11 months ago

guys this question actually came out on the actual test!

upvoted 7 times

🗨️ 👤 **smishra** 4 years, 11 months ago

We don't have "true" and "false" filters as an option available, why so? What are we supposed to answer for future orders (false filter) and all order (true filter)

upvoted 1 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:

**Suggested Answer:** *Explanation*

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.

References:

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

 **Daltonic75** Highly Voted 5 years, 2 months ago

Options: <https://www.examttopics.com/assets/media/exam-media/02521/0004000001.jpg>  
upvoted 47 times

 **meoukg** Most Recent 3 years, 2 months ago

Got it on 03/2022, chose the same as the given answer.  
upvoted 1 times

 **and\_p01** 4 years, 5 months ago

Correct answer.  
upvoted 4 times

You develop a solution that uses an Azure SQL Database to store user information for a mobile app.

The app stores sensitive information about users.

You need to hide sensitive information from developers that query the data for the mobile app.

Which three items must you identify when configuring dynamic data masking? Each correct answer presents a part of the solution.

NOTE: Each correct selection is worth one point.

- A. Column
- B. Table
- C. Trigger
- D. Index
- E. Schema

**Suggested Answer: ABE**

In the Dynamic Data Masking configuration page, you may see some database columns that the recommendations engine has flagged for masking. In order to accept the recommendations, just click Add Mask for one or more columns and a mask is created based on the default type for this column. You can change the masking function by clicking on the masking rule and editing the masking field format to a different format of your choice.

SCHEMA	TABLE	COLUMN	
SalesLT	Customer	FirstName	ADD MASK
SalesLT	Customer	LastName	ADD MASK
SalesLT	Customer	EmailAddress	ADD MASK
SalesLT	Customer	Phone	ADD MASK
SalesLT	CustomerAddress	AddressID	ADD MASK

References:

<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-dynamic-data-masking-get-started-portal>

**SidSadhu** Highly Voted 4 years, 9 months ago

Correct

upvoted 7 times

**GCMAN** Most Recent 4 years, 6 months ago

GAC = given answer correct

upvoted 3 times

## HOTSPOT -

You store customer information in an 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. (Line numbers are included for reference only.)

```

1 CloudTableClient tableClient = account.CreateCloudTableClient();
2 CloudTable table = tableClient.GetTableReference ("people");
3 TableQuery<CustomerEntity> query = new TableQuery < CustomerEntity > ()
4.     .Where (TableQuery.CombineFilters (
5     TableQuery.GenerateFilterCondition (PartitionKey, QueryComparisons.Equal, "Smith"),
6     TableOperators.And, TableQuery.
7     GenerateFilterCondition (Email, QueryComparisons.Equal, "ssmith@contoso.com")
8 );
8 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

Yes

No

The code returns every Record where the surname equals **Smith**.

☐
☐

The table endpoint `https://<mytableendpoint>/People`

(PartitionKey='Smith',RowKey='Steve') returns the same results as the code.

☐
☐

Suggested Answer:

## Answer Area

Yes

No

The code returns every Record where the surname equals **Smith**.

☐
☒

The table endpoint `https://<mytableendpoint>/People`

(PartitionKey='Smith',RowKey='Steve') returns the same results as the code.

☒
☐

Box 1: No -

Only the second row is returned thank to the Email filter condition.

Box 2: Yes -

This also returns the second row.

 **GCMan** Highly Voted 4 years, 6 months ago

correct

upvoted 7 times

A company uses Azure SQL Database to store data for an app. The data includes sensitive information.

You need to implement measures that allow only members of the managers group to see sensitive information.

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

NOTE: Each correct selection is worth one point.

A. Include the managers group.

B. Exclude the managers group.

C. Exclude the administrators group.

D. Navigate to the following URL:

```
PUT https://management.azure.com/subscriptions/00000000-1111-2222-3333-444444444444
/resourceGroups/rg01/providers/Microsoft.Sql/servers/server01/databases/customers
/transparentDataEncryption/current?api-version=2014-04-01
```

E. Run the following Azure PowerShell command:

```
New-AzureRmSqlDatabaseDataMaskingRule -SchemaName "dbo" -TableName "customers" `
-ColumnName "ssn" -MaskingFunction "Default"
```

#### Suggested Answer: BE

Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer.

SQL users excluded from masking - A set of SQL users or AAD identities that get unmasked data in the SQL query results.

Note: The New-AzureRmSqlDatabaseDataMaskingRule cmdlet creates a data masking rule for an Azure SQL database.

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.sql/new-azurermSqlDatabaseDataMaskingRule?view=azurerm-ps-6.13.0>

 **saeza100** Highly Voted 4 years, 11 months ago

B and E are correct answers

upvoted 24 times


 **nikos2001** Highly Voted 4 years, 11 months ago

B,E

See "SQL users excluded from masking"


<https://docs.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-configure-portal>

upvoted 16 times

 **LTiwana** 4 years, 10 months ago

Thanks for the link it cleared my confusion. Answer is B and E

upvoted 3 times

 **NaliniTomar** Most Recent 4 years, 9 months ago

Given answer is correct.

upvoted 3 times

 **navaneeth** 4 years, 9 months ago

if we exclude how come we can achieve this.No wrong answer please.implement measures that allow only members of the managers group to see sensitive information

upvoted 1 times

 **Cornholioz** 4 years, 5 months ago

E adds a mask to the SSN column. By "excluding" the Managers from this masking (B), the column gets masked for everyone other than Managers.

upvoted 2 times

 **dkaps06** 4 years, 9 months ago

Yup agree with given answer. B and E are correct choices

upvoted 4 times

 **chaudh** 4 years, 10 months ago



BE is correct

upvoted 4 times

🗳️ 👤 **brajen** 4 years, 10 months ago

Given Answers Is Correct Only --BE

upvoted 6 times

🗳️ 👤 **superman917** 4 years, 11 months ago

Might be "Exclude Managers Group".

<https://www.sqlshack.com/implementing-dynamic-data-masking-in-azure-sql-database/>

upvoted 1 times

🗳️ 👤 **AK89** 5 years ago

Correct Answer are A and E

upvoted 3 times

🗳️ 👤 **alirasouli** 5 years ago

"Included" means masking is disabled for that group/user. Public users (means not logged in users) are excluded. Here A & E are the correct choices.

upvoted 3 times

🗳️ 👤 **Berlin** 5 years, 2 months ago

Question is to allow access to managers group for sensitive information. If we Exclude Managers Group(option B), How we can enforce the sensitivity.

upvoted 4 times

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

We are data masking so all users will not be able to see the sensitive information. Only manager will be able to see if you exclude them.

upvoted 16 times

🗳️ 👤 **rrongcheng** 4 years, 10 months ago

By -MaskingFunction 'Default'

Masking logic: Full masking according to the data types of the designated fields

<https://docs.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-overview>

upvoted 2 times

🗳️ 👤 **kirannep** 4 years, 7 months ago

Question is not valid. It doesn't say about masking only the options have abstract meaning of it.

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 in network outages or other unforeseen failures
- ⇒ Process reservations in the exact sequence as reservation 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:

**Suggested Answer:** *Explanation*

Box 1: 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.

" Process reservations in the exact sequence as reservation are submitted to minimize overbooking or selling the same seat to multiple travelers."

Box 2: Enable-automatic-failover

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.

" Accept reservations even in network outages or other unforeseen failures"

Box 3: southcentralus -

" You provision a resource group named `airlineResourceGroup` in the Azure South-Central US region."

References:

<https://docs.microsoft.com/en-us/azure/cosmos-db/consistency-levels> <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/cosmos-db/high-availability.md>

AK89 Highly Voted 5 years ago

The correct answer in sequence is:

Bounded Staleness

--enable-automatic-failover true

'southcentralus=0 eastus=1 westus-2'

upvoted 63 times

TaiH Highly Voted 5 years, 4 months ago

Based on this:

- ⇒ Process reservations in the exact sequence as reservation 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.

The first box should be staleness.

upvoted 54 times

meoukg Most Recent 3 years, 2 months ago

Got it on 03/2022, I chose as below (not sure it was correct :)))

1. Strong

2. --enable-automatic-failover true

3. --locations'southcentralus=0 eastus=1 westus=2

upvoted 1 times

dancsita 4 years, 3 months ago

According to the docs:

--enable-automatic-failover

Enables automatic failover of the write region in the rare event that the region is unavailable due to an outage. Automatic failover will result in a new write region for the account and is chosen based on the failover priorities configured for the account.

upvoted 1 times

🗳️ 👤 **Cornholioz** 4 years, 5 months ago

Tried the CLI in the portal's Cloud Shell:

--locations has wrong syntax in the options.

```
az cosmosdb create -n myaccount4 --enable-automatic-failover true --max-interval 5 -g myRG --locations regionname=eastus --default-consistency-level=BoundedStaleness
```

--locations needs regionname. Because it also takes in failoverPriority and isZoneRedundant params.

So either the option with just the SouthCentralUS should work or if the rest of the regions need to be mentioned, then it's a deprecated syntax which makes no sense to answer now.

upvoted 1 times

🗳️ 👤 **Nitin121** 4 years, 7 months ago

Consistency level should be Bounded Staleness

--max-interval

When used with Bounded Staleness consistency, this value represents the time amount of staleness (in seconds) tolerated. Accepted range for this value is 1 - 100.

default value: 5

[https://docs.microsoft.com/en-us/cli/azure/cosmosdb?view=azure-cli-latest#az\\_cosmosdb\\_create](https://docs.microsoft.com/en-us/cli/azure/cosmosdb?view=azure-cli-latest#az_cosmosdb_create)

upvoted 1 times

🗳️ 👤 **Sumeet\_das** 4 years, 7 months ago

Same question is there in AZ204, correct answer for 1st is Bounded staleness. Other 2 is fine

upvoted 1 times

🗳️ 👤 **NPoul\_92** 4 years, 8 months ago

<https://vceguide.com/how-should-you-complete-the-azure-cli-commands/>

upvoted 1 times

🗳️ 👤 **malawneh** 4 years, 9 months ago

I passed today with score 860, i got this in my exam.

My answer was:

Bounded Staleness

--enable-automatic-failover true

'southcentralus=0 eastus=1 westus=2'

upvoted 9 times

🗳️ 👤 **sinh** 4 years, 7 months ago

You may not be the right answer.

upvoted 2 times

🗳️ 👤 **chaudh** 4 years, 10 months ago

1. BoundedStaleness: question doesn't require 100% consistency (only exact sequence, minimize overbooking or selling), and also allow for 5-second tolerance window.

2. --enable-automatic-failover

3. --locations 'southcentralus'. Note that 'southcentralus=0...' is the failover policy, not locations as specified.

upvoted 9 times

🗳️ 👤 **Satheesh88** 4 years, 5 months ago

[https://docs.microsoft.com/en-us/azure/cosmos-db/high-](https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability#:~:text=For%20multi%2Dregion%20Azure%20Cosmos,will%20automatically%20failover%20your%20account.)

[availability#:~:text=For%20multi%2Dregion%20Azure%20Cosmos,will%20automatically%20failover%20your%20account.](https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability#:~:text=For%20multi%2Dregion%20Azure%20Cosmos,will%20automatically%20failover%20your%20account.)

upvoted 1 times

🗳️ 👤 **Satheesh88** 4 years, 5 months ago

Multiple write region is set by fail over policies.



az cosmosdb failover-priority-change --ids \$accountId \

--failover-policies 'West US 2=0' 'South Central US=1' 'East US 2=2'

Ref: <https://docs.microsoft.com/en-us/azure/cosmos-db/manage-with-cli#enable-automatic-failover>

Box 3 - SouthcentralUs

upvoted 2 times

  **Juanlu** 4 years, 3 months ago



Thanks, they look like good arguments, so:

Box 1: BoundedStaleness

Box2: --enable-automatic-failover

Box 3: SouthCentra US

upvoted 1 times

  **Raevex** 4 years, 10 months ago

Options:

[https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020\\_Page\\_139\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020_Page_139_Image_0001.jpg)

upvoted 16 times

  **Suchismita** 4 years, 11 months ago

The correct answer in sequence is:

Bounded Staleness

--enable-automatic-failover true

southcentralus - As provision is to create a resource group named airlineResourceGroup in the Azure South-Central US region.

upvoted 1 times

  **Dayakar** 5 years ago



In the question it is mentioned that low latency , In this case it should be bounded staleness. If you want high latency then only strong consistency

upvoted 3 times

  **Daltonic75** 5 years, 2 months ago

Options: <https://www.sharetest.com/images/AZ-202.files/image095.jpg>

upvoted 27 times

  **Khang** 5 years, 3 months ago

First box should be: Staleness

Allow simultaneous and out-of-order reservations with a maximum five-second tolerance window.



upvoted 19 times

  **boldarev** 5 years, 3 months ago

☞ Accept reservations even in network outages or other unforeseen failures

Box 3 should be with several locations

upvoted 37 times

  **LTiwana** 4 years, 10 months ago

I think you are right. I found 2 links that explain this:

<https://docs.azure.cn/zh-cn/cli/cosmosdb?view=azure-cli-latest>

--locations

Space-separated locations in 'regionName=failoverPriority' format. E.g eastus=0 westus=1.

<https://stackoverflow.com/questions/51197375/creating-cosmosdb-with-azure-cli-in-powershell>

upvoted 3 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 in real time when specific requirements are met, excluding backup blob copies. What should you do?

- A. Download the blob to a virtual machine and then upload the blob to Container2.
- B. Run the Azure PowerShell command Start-AzureStorageBlobCopy.
- C. Copy blobs to Container2 by using the Put Blob operation of the Blob Service REST API.
- D. Use AzCopy with the Snapshot switch blobs 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.

References:

<https://docs.microsoft.com/en-us/powershell/module/azure.storage/start-azurestorageblobcopy?view=azurermps-6.13.0>

  **chaudh**  4 years, 10 months ago

why not C?

The REST API for uploading is already supported. Also, coping with PowerShell or CLI seems not to be good choices for "real time" requirement.

upvoted 12 times

  **Bluediamond** 4 years, 10 months ago

I tend to agree with you. Now if B said something like create a PS runbook that is set on a webhook connected to storage account then that would make sense.

upvoted 3 times

  **rrongcheng** 4 years, 10 months ago

I think it's C too. As it uses storage REST API to upload to Container1, it's easy to do the same to put into Container2 at the same time to make it 'in real time'

upvoted 3 times

  **altafpatel1984**  3 years, 6 months ago

If Start-AzureStorageBlobCopy is triggered by Event Grid then it make sense, which is not mentioned here so C is correct.

For following question, since Event Grid is mentioned with Start-AzureStorageBlobCopy, that make sense:

<https://www.examttopics.com/discussions/microsoft/view/37714-exam-az-204-topic-3-question-24-discussion/>



upvoted 2 times

  **Mehabooba** 4 years, 4 months ago

The Put Blob operation creates a new block, page, or append blob, or updates the content of an existing block blob. Updating an existing block blob overwrites any existing metadata on the blob.



The difference between the two is likely AzCopy is operating a bulk mode and Start-AzureStorageBlobCopy is operating serially due to piping the blobs in one by one, sending request to initiate the transfer

upvoted 1 times


  **Chilred** 4 years, 4 months ago

There is a similar question in az204: <https://www.examttopics.com/discussions/microsoft/view/37714-exam-az-204-topic-3-question-7-discussion/>

In both the option C is available. So in choice between B and C, I would go with C  
upvoted 1 times

  **Chilred** 4 years, 4 months ago



Sorry, my fault. This are different questions....  
upvoted 4 times

  **kirannep** 4 years, 7 months ago

Anyone saying put, The Put Blob operation creates a new block, page, or append blob, or updates the content of an existing block blob. Answer is B because you want to copy in realtime even if it is ps script.  
upvoted 3 times

  **Larry616** 4 years, 9 months ago



Answer should be C  
B only makes sense if it was "Create a PS script and set it up by AZ Functions or Webjobs", but it's "Run PS ..." which means triggered manually and couldn't meet the "real time" requirement  
upvoted 2 times

  **xsify** 4 years, 10 months ago

Copy with Put Blob (C) requires more step, for example, Put Blob on page blob and append blob only initialize the blob, additional step is needed to copy the blob. So i think the answers is B.  
<https://docs.microsoft.com/en-us/rest/api/storageservices/put-blob>  
upvoted 4 times

  **rrongcheng** 4 years, 10 months ago

B doesn't meet the 'real-time' requirement. It has to be done by code.  
upvoted 1 times

  **Kobee** 4 years, 5 months ago

It uses a storage REST API to upload media, so the block type is a Block Blob, not a Page Blob (used for VM) and Append Blob (used to append data like logs)  
So I think the correct answer is C.  
For a block blob, the request body contains the content of the blob (see <https://docs.microsoft.com/en-us/rest/api/storageservices/put-blob#request-body>)  
upvoted 1 times

## HOTSPOT -

You have an app that stores player scores for an online game. The app stores data in Azure tables using a class named PlayerScore as the table entity. The table is populated with 100,000 records.

You are reviewing the following section of code that is intended to retrieve 20 records where the player score exceeds 15,000. (Line numbers are included for reference only.)

```
1 public void GetScore(string playerId, int score, string gameName)
2 {
3     TableQuery<DynamicTableEntity> query = new TableQuery<DynamicTableEntity>().Select(new string[] { "Score" })
4         .Where(TableQuery.GenerateFilterConditionForInt("Score", QueryComparisons.GreaterThanOrEqualTo, 15000)).Take
5         (20);
6     EntityResolver<KeyValuePair<string, int?>> resolver =
7         (partitionKey, rowKey, ts, props, etag) => new KeyValuePair<string, int?>(rowKey, props["Score"].Int32Value);
8     foreach (var scoreItem in scoreTable.ExecuteQuery(query, resolver, null, null))
9     {
10         Console.WriteLine($"{scoreItem.Key} {scoreItem.Value}");
11     }
12 }

13 public class PlayerScore : TableEntity
14 {
15     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
16     {
17         PartitionKey = gameId;
18         RowKey = playerId;
19         Score = score;
20         TimePlayed = timePlayed;
21     }
22     public int Score { get; set; }
23     public long TimePlayed { get; set; }
24 }
```

You have the following code. (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 }
11 public class PlayerScore : TableEntity
12 {
13     public PlayerScore(string gameId, string playerId, int score, long timePlayed)
14     {
15         this.PartitionKey = gameId;
16         this.RowKey = playerId;
17         Score = score;
18         TimePlayed = timePlayed;
19     }
20     public int Score { get; set; }
21     public long TimePlayed { get; set; }
22 }
```

You store customer information in an Azure Cosmos database. 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. (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.GenerateAnd, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal, "Smith"),
06         TableOperators.And, TableQuery.GenerateFilterCondition(Email, QueryComparisons.Equal,
07             "ssmith@contoso.com")
08     ));
09 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:

**Suggested Answer:** *Explanation*

Box 1: No -

Box 2: Yes -

The TableQuery.Take method defines the upper bound for the number of entities the query returns.

Example:

```
query.Take(10);
```

Box 3: Yes -

Box 4: Yes -

References:

<https://www.vkinfotek.com/azureqa/how-do-i-query-azure-table-storage-using-tablequery-class.html>

🗨️ **Daltonic75** Highly Voted 5 years, 2 months ago

[https://www.itexams.com/static/img/exams/Microsoft-AZ-203-1.0/xmlfile-177\\_1.png](https://www.itexams.com/static/img/exams/Microsoft-AZ-203-1.0/xmlfile-177_1.png)

upvoted 28 times

🗨️ **chaudh** 4 years, 10 months ago

[https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020\\_Page\\_143\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020_Page_143_Image_0001.jpg)

upvoted 19 times

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

Daltonic75, I don't know who you are. But clearly, examtopics should pay you for every subscriptions they receive as you did 50% of their work. Thank you.

upvoted 68 times

🗨️ **bhushan\_786** 4 years, 6 months ago

The link provided by him is not relevant here :)

It is meant for other question.

upvoted 6 times

🗨️ **altafpatel1984** 3 years, 5 months ago

but this time he did wrong job ! posted different image.

upvoted 1 times

🗨️ **Ric0720** Highly Voted 4 years, 9 months ago

No,Yes,No,Yes

upvoted 15 times

🗨️ **Juanlu** Most Recent 4 years, 3 months ago

NO

YES.

YES. All record are returned to the client. The maximum value are 1000. If so, tokencontinuation returned and we can use:

ExecuteQuerySegmentedAsync

YES. By default PartitionKey and RowKey are already included by default in the select clause.

<https://www.examttopics.com/discussions/microsoft/view/16363-exam-az-203-topic-24-question-13-discussion/>

upvoted 1 times

🗨️ **Juanlu** 4 years, 3 months ago

In my opinion, all answer are correct as the solution mention, so: YES, NO, NO, NO !

upvoted 1 times

🗨️ **NPaul\_92** 4 years, 8 months ago

<https://vceguide.com/where-the-player-score-exceeds-15000-line-numbers-are-included-for-reference-only/>

upvoted 2 times

🗨️ **nikos2001** 4 years, 11 months ago

Third box should be NO. The client will get not all records, but those fitting the query filter.

upvoted 12 times

🗨️ **godsbane** 4 years, 10 months ago



Box 3 is No. If .Take() is not specified, a maximum of 1000 records are returned to the client. There is no way all records are returned to the client, what if we have 20 mil records?

Reference: <https://docs.microsoft.com/bs-latn-ba/java/api/com.microsoft.azure.storage.table.tablequery.take?view=azure-java-legacy>

upvoted 7 times

🗨️ 👤 **Praveen24by7** 4 years, 11 months ago

The code queries the Azure table and retrieves the TimePlayed property from the table - No

The code will display a maximum of twenty records - Yes

All records will be sent to the client. The client will display records for scores greater than or equal to 15,000. - Yes

The scoreItem.Key property of the KeyValuePairs that ExecuteQuery returns will contain a value for PlayerID. - Yes

upvoted 4 times

🗨️ 👤 **pranay** 4 years, 10 months ago

Third answer should be NO

upvoted 5 times

🗨️ 👤 **Chilred** 4 years, 5 months ago

Then explain why

upvoted 1 times

🗨️ 👤 **saeza100** 4 years, 11 months ago

Box4 should be "No" , we select just "Score" with "Query"

upvoted 1 times

🗨️ 👤 **pranay** 4 years, 10 months ago

PlayerId is a rowKey so it will be returned anyhow. You don't need to mention rowKey and PartitionKey.

upvoted 1 times

🗨️ 👤 **Juanlu** 4 years, 3 months ago

yes, but as I know, Partitionkey and Rowkey are also included by default, so de answer should be YES.

upvoted 1 times

🗨️ 👤 **lau13** 5 years ago

options

[Yes|No]The code queries the Azure table and retrieves the TimePlayed property from the table

[Yes|No]The code will display a maximum of twenty records

[Yes|No]All records will be sent to the client. The client will display records for scores greater than or equal to 15,000.

[Yes|No]The scoreItem.Key property of the KeyValuePairs that ExecuteQuery returns will contain a value for PlayerID.

upvoted 6 times

🗨️ 👤 **realbart** 5 years, 1 month ago

Not all records, but only 20 records will not be sent to the client but written to the console

upvoted 4 times

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

Box 3: should be false, please suggest

upvoted 8 times

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

Why do you think so? I think it is correct.

upvoted 1 times

🗨️ 👤 **LEOPOLD** 5 years, 1 month ago

Same conclusion. I thought it would be deferred execution and iteration over the query with execute it so you have a query with .take 20 in it and when iterating you dont get all records sent to the client. (I sure hope thats what happens. Imagine that you had 2000000000 records in that DB all sent to client :D)

upvoted 5 times

🗨️ 👤 **khrystyna** 5 years, 1 month ago



It can be because of method order:

Select - get all records to client

Where - in memory filtering on client side



Take - return just 20 records

upvoted 2 times

  **Ummara** 4 years, 9 months ago

the partition key is the game id, and select is not on the partition key. In such a case all records are retrieved, called Table Scan, and therefore Client will filter records from the data fetched.

upvoted 2 times

  **Ummara** 4 years, 9 months ago

A Table Scan does not include the PartitionKey and is very inefficient because it searches all of the partitions that make up your table in turn for any matching entities. It will perform a table scan regardless of whether or not your filter uses the RowKey. For example: \$filter=LastName eq 'Jones'

<https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-design-for-query>

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. Determine whether the solution meets the stated goals.

You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

- ⇒ Share session state across all ASP.NET web applications
- ⇒ Support controlled, concurrent access to the same session state data for multiple readers and a single writer
- ⇒ Save full HTTP responses for concurrent requests

You need to store the information.

Proposed Solution: Deploy and configure an Azure Database for PostgreSQL. Update the web applications.

Does the solution meet the goal?

A. Yes


B. No

**Suggested Answer: B**

  **Drunken**  4 years, 11 months ago

Redis Cache is better for this but Posgre would also work (although it performs in this case worse regarding scalability). However the question is misleading since it wants you to select B (No) but it would be a valid choice although inefficient...

upvoted 18 times

  **pm\_ufa** 4 years, 5 months ago



I totally agree, PgSQL by itself is a viable option, but i think extra conditions like `access to the same session state data for multiple readers and a single writer` and `Save full HTTP responses for concurrent requests` are meant to invalidate it

upvoted 3 times

  **bhushan\_786** 4 years, 5 months ago

So what should be answer to it? Yes or No??

upvoted 1 times

  **m\_siri** 4 years, 5 months ago

the provided answer is correct

ref: <https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching#cached-session-state-and-html-output>

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. Determine whether the solution meets the stated goals.

You are developing and deploying several ASP.Net web applications to Azure App Service. You plan to save session state information and HTML output. You must use a storage mechanism with the following requirements:

- ⇒ Share session state across all ASP.NET web applications
- ⇒ Support controlled, concurrent access to the same session state data for multiple readers and a single writer
- ⇒ Save full HTTP responses for concurrent requests

You need to store the information.

Proposed Solution: Deploy and configure Azure Cache for Redis. Update the web applications.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

  **EYIT**  5 years, 6 months ago


Wrong key - it should be YES: <https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching?source=docs#cached-session-state-and-html-output>  
upvoted 65 times

  **gclopes**  5 years, 2 months ago



One more vote for the A answer!!  
upvoted 17 times

  **ss22**  4 years, 3 months ago



Correct Answer- Yes  
upvoted 2 times

  **goudigubba** 4 years, 3 months ago

Its a simple question yet the answer is given as wrong. No point of providing wrong answers.  
upvoted 1 times

  **bashirk** 4 years, 9 months ago

Correct Ans. Yes  
upvoted 2 times

  **hertino** 4 years, 9 months ago

Today I got 900, I had this question, my answer: Yes  
upvoted 7 times

  **Anamika7** 4 years, 9 months ago

Hey hertino, if you have the list of all the question with answer. Could you please share it  
upvoted 14 times

  **mcampos** 4 years, 11 months ago

YES  
<https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching>  
upvoted 1 times

  **AK89** 5 years ago

It is YES  
upvoted 2 times



  **Dayakar** 5 years ago

It should be "Yes", because using redis cache store the session information and use it for later point of time.  
upvoted 5 times

  **dollarpo7** 5 years ago



Answer is Yes

upvoted 4 times

  **heero** 5 years, 3 months ago

Answer: A

upvoted 12 times

  **Khang** 5 years, 4 months ago

Correct answer is A (Yes)

upvoted 11 times

  **shanky\_007** 5 years, 5 months ago

correct Ans: Yes

upvoted 6 times

  **mohanadzaidi** 5 years, 6 months ago

Should be Yes ! thanks EYIT

upvoted 13 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 question, 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 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 Hub. Configure the machine identifier as the partition key and enable capture.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: A**

References:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide>

  **Larry616** Highly Voted 4 years, 9 months ago



Answer should be B-No

Because total partitions has a max# of 32, but in the question it uses machine identifier as partition key, there're  $2000 * 5 = 10,000$  machines here

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features#partitions>

The number of partitions is specified at creation and must be between 2 and 32

upvoted 12 times

  **Brak** 4 years, 3 months ago

Wrong. See my response above.

upvoted 1 times

  **clarionprogrammer** 4 years, 1 month ago

This is wrong. You're confusing Partition Keys with creating partitions.

upvoted 4 times

  **Secure01** Highly Voted 4 years, 6 months ago

1. Partition Key is not mapped 1:1 to partition. 2. Store not accumulate data per each device. Each device send data separately. 3. 2MB per 24 hours are not sent at once. 4. Event Hub supposed to take events/messages from big amount of devices. 5. Event still a message (small one). You can put information you need to store in Blob. 6. To store message in Blob with Event Hub is easy. Just switch on capture setting.

upvoted 8 times

  **Brak** 4 years, 3 months ago

Correct. Event Hub + Capture is a no-code solution that can scale.

upvoted 1 times

  **meoukg** Most Recent 3 years, 2 months ago

Got it on 03/2022, I chose A. Yes

upvoted 2 times



  **RahulKate** 4 years, 4 months ago

Partitions - 2000 per CU - Dedicated Plan - So we can have 2K+ partitions in Event Hub

Size allowed - Yes

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-dedicated-overview#event-hubs-dedicated-quotas-and-limits>

upvoted 2 times

  **Esward** 2 years, 3 months ago

Yes is the correct answer

<https://learn.microsoft.com/en-us/azure/event-hubs/event-hubs-dedicated-overview#event-hubs-dedicated-quotas-and-limits>

upvoted 1 times

  **ExamStudent123** 4 years, 9 months ago

It is not Event Hub for the following reasons:



- "The event data has information about what happened but doesn't have the data that triggered the event." <https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services>
- "Maximum size of Event Hubs event: 1 MB (standard)". From the question "A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data." That means the max size from one store could be 2MB to 10MB.

Service Bus is correct for the following reasons:

- "A message is raw data produced by a service to be consumed or stored elsewhere."
- "Order processing and financial transactions"

<https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services>



upvoted 7 times

  **godsbane** 4 years, 10 months ago

Answer should be NO.

<https://azure.microsoft.com/es-es/blog/events-data-points-and-messages-choosing-the-right-azure-messaging-service-for-your-data/>



upvoted 5 times

  **LTiwana** 4 years, 10 months ago

In this like MS is talking about using device identity for partition key under Event Hub:



The temperature sensors discussed earlier each emit a distinct stream that will be kept together using the Event Hub partitioning model, using the identity of the device as the partitioning key.

upvoted 1 times

  **nickbilak** 4 years, 11 months ago



but, if we should transfer each sale from POS (not the entire daily batch at once), I'd go with Service Bus

upvoted 1 times

  **nickbilak** 4 years, 11 months ago



scratch that, eventdatabatch.tryadd is not for service bus, but for event hub. SB has 256K limit. then the right solution could be event grid triggered event for consumers upon blob upload, that event will have id/link to the blob with 2MB of POS data. Event hub can have only up to 32 partitions, so answer A has no sense.

upvoted 2 times

  **Brak** 4 years, 3 months ago

32 is the limit for physical partitions. Each device id is hashed to choose a physical partition.



upvoted 1 times

  **nickbilak** 4 years, 11 months ago

I think Service Bus is correct as it's used for financial transactions (with guaranteed delivery) and allows batches up to 4MB -



<https://docs.microsoft.com/en-us/dotnet/api/microsoft.servicebus.messaging.eventdatabatch.tryadd?view=azure-dotnet#remarks>

upvoted 5 times

  **nickbilak** 4 years, 11 months ago

The event hub message size limit is 1MB. How can be answer A be correct, if the requirement is 2MB of data?

upvoted 2 times

  **LTiwana** 4 years, 10 months ago

But the question mentions 2MB data every 24 hours and doesn't specify if the data needs to be batched or streamed (Event Hub accepts both).

It's confusion, but the answer might be correct. Here is a good comparison between the three services:

<https://docs.microsoft.com/en-us/azure/event-grid/compare-messaging-services>

upvoted 6 times

  **rrongcheng** 4 years, 10 months ago

I think it's Yes.

In this linked doc, it gives an example at the end. The example uses Event hubs to 'Capture file' to Storage, and then 'Capture event' to next step.

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 question, 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 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 event filtering to evaluate the device identifier.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

References:

<https://docs.microsoft.com/en-us/azure/event-grid/event-filtering>

  **GraceCyborg** 4 years, 5 months ago

service bus

upvoted 1 times

  **subrata83** 4 years, 8 months ago

Answer seems correct. How can Event grid store the data. Here the requirement is to store the blob

upvoted 1 times

  **Dayakar** 5 years ago

Azure event grid supports upto 1 mega byte only but in the question it is mentioned 2 mega bytes hence azure event grid is not the right answer.

upvoted 2 times

  **gayan1991** 4 years, 10 months ago

What about the part where it is mentioned that 2 MB for 24 hours?



upvoted 2 times

  **ExamStudent123** 4 years, 9 months ago

"A single device can produce 2 megabytes (MB) of data every 24 hours. Each store location has one to five devices that send data."

It is a for a single device, and a store can have up-to 5 devices. So the max payload could be 2MB to 10MB.

upvoted 1 times

  **mohanadzaidi** 5 years, 6 months ago

Is not the Answer "Yes": EventGrid can manage the destination resource depending on the source of the event.

upvoted 2 times

  **motu** 5 years, 3 months ago

Event Grid only supports messages up to 1 MB, 2 MB is asked for in the question

upvoted 9 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 question, 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 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 Notification Hub. Register all devices with the hub.

Does the solution meet the goal?

A. Yes

B. No

**Suggested Answer: B**

Instead provision an Azure Event Hub. Configure the machine identifier as the partition key and enable capture.

References:

<https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-programming-guide>

  **Eysi** 4 years, 6 months ago



Could somebody explain why Azure Notification Hub is not appropriate solution? Thanks.

upvoted 2 times

  **m\_siri** 4 years, 5 months ago

Azure Notification Hub is for only to push data to the client devices but not to receive data from client devices.

upvoted 3 times

  **BizTalk** 4 years, 9 months ago

It is correct

upvoted 1 times

You develop and deploy a Java RESTful API to Azure App Service.

You open a browser and navigate to the URL for the API. You receive the following error message:

```
Failed to load http://api.azurewebsites.net:6000/#/api/Products: No 'Access-Control-Allow-Origin' header is present on the requested resource.
Origin 'http://localhost:6000' is therefore not allowed access
```

You need to resolve the error.

What should you do?

- A. Bind an SSL certificate
- B. Enable authentication
- C. Enable CORS
- D. Map a custom domain
- E. Add a CDN

**Suggested Answer: C**

We need to enable Cross-Origin Resource Sharing (CORS).

References:

<https://medium.com/@xinganwang/a-practical-guide-to-cors-51e8fd329a1f>

 **bashirk**  4 years, 9 months ago

Answer is correct

upvoted 5 times

## HOTSPOT -

Your company is migrating applications to Azure. The IT department must allow internal developers to communicate with Microsoft support. The service agents of the IT department must only have view resources and create support ticket permissions to all subscriptions. A new custom role must be created by reusing a default role definition and changing the permissions.

You need to create the custom role.

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Suggested Answer:** *Explanation*

Box 1: Set-AzureRmRoleDefinition Input-File C:\SupportRole.json

The Set-AzureRmRoleDefinition cmdlet updates an existing custom role in Azure Role-Based Access Control. Provide the updated role definition as an input to the command as a JSON file or a PSRoleDefinition object.

The role definition for the updated custom role MUST contain the Id and all other required properties of the role even if they are not updated:

DisplayName,

Description, Actions, AssignableScope

Box 2: "\*/read\*.\* Microsoft.Support/\*"

Microsoft.Support/\* Create and manage support tickets

"Microsoft.Support" role definition azure

Incorrect Answers:




Get-AzureRmRoleDefinition. The Get-AzureRmRoleDefinition command does not have an action section.

First, use the Get-AzureRmRoleDefinition command to retrieve the custom role that you wish to modify. Then, modify the properties that you wish to change.

Finally, save the role definition using the Set-AzureRmRoleDefinition command.

References:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/custom-roles-powershell>

  **AK89**  5 years ago

Box1: Get-AzRoleDefinition -Name "Reader" | ConvertTo-Json Out-File C:\Sample.json

Box2: Set-AzureRmRoleDefinition Input-File C:\SupportRole.json

Box3: "read/\*,Microsoft.Support/\*"

upvoted 35 times

  **Larry616** 4 years, 9 months ago

Why only read? The question mentioned it also need "create support ticket permissions"

upvoted 1 times

  **Larry616** 4 years, 9 months ago

Got it now, box3 is suffix of command in box2, it appends "microsoft.support" permission to existing "read" permission, above answer is correct

upvoted 2 times

  **t8203**  4 years, 9 months ago

options: [https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020\\_Page\\_183\\_Image\\_0001.jpg](https://vceguide.com/wp-content/uploads/2020/05/Microsoft-AZ-203-date-05-08-2020_Page_183_Image_0001.jpg)

upvoted 14 times

  **Sylph**  4 years, 2 months ago

Tutorial for this exact situation: <https://docs.microsoft.com/en-us/azure/role-based-access-control/tutorial-custom-role-powershell>

upvoted 2 times

  **VK\_Gladiator** 4 years, 4 months ago

Given answer is correct.

Actions Description

\*/read Read resources of all types, except secrets.

Microsoft.Authorization/\* Manage authorization

Microsoft.Support/\* Create and update a support ticket

Reference link: <https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

upvoted 1 times

🗨️ 👤 **AnkitN** 4 years, 9 months ago

If you look at the built-in role and see an example for actions to allow support to raise Microsoft support tickets, you can see the required action:

Since this is clearly shown in the documentation, all other options are incorrect.

For more information on the built-in roles, one can go to the below link:

<https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>

upvoted 1 times

🗨️ 👤 **Steve\_az** 4 years, 9 months ago

<https://vceguide.com/hotspot-905/>

upvoted 3 times

🗨️ 👤 **sjsouvik** 5 years ago

answer should be :

Box1: Get-AzRoleDefinition -Name "Reader" | ConvertTo-Json Out-File C:\Sample.json

Box2: Set-AzureRmRoleDefinition Input-File C:\SupportRole.json

Box3: "Microsoft.Support/\*"

upvoted 5 times

You are developing an internal website for employees to view sensitive data. The website uses Azure Active Directory (AAD) for authentication. You need to implement multifactor authentication for the website. What should you do? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Upgrade to Azure AD Premium.
- B. In Azure AD conditional access, enable the baseline policy.
- C. In Azure AD, create a new conditional access policy.
- D. In Azure AD, enable application proxy.
- E. Configure the website to use Azure AD B2C.

**Suggested Answer: AC**

A: Multi-Factor Authentication comes as part of the following offerings:

- ⇒ Azure Active Directory Premium licenses - Full featured use of Azure Multi-Factor Authentication Service (Cloud) or Azure Multi-Factor Authentication Server (On-premises).
- ⇒ Multi-Factor Authentication for Office 365
- ⇒ Azure Active Directory Global Administrators

C: MFA Enabled by conditional access policy. It is the most flexible means to enable two-step verification for your users. Enabling using conditional access policy only works for Azure MFA in the cloud and is a premium feature of Azure AD.

References:

<https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-getstarted>

 **Larry616** Highly Voted 4 years, 9 months ago

Answer is correct

upvoted 9 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:

**Suggested Answer:** *Explanation*

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)" />
Etc.
```

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>

References:

<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>

  **Sam1101**  4 years, 11 months ago

1. Inbound
2. Inbound
3. Inbound
4. Outbound

upvoted 33 times

  **TaiH**  5 years, 4 months ago

#3: cache-store value should be inbound.

Ref: <https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

upvoted 19 times

  **LEOPOLD** 5 years, 1 month ago

Hold up in that link it says:

Usage

This policy can be used in the following policy sections and scopes.

Policy sections: inbound, outbound, backend, on-error

Policy scopes: all scopes

upvoted 2 times

  **LEOPOLD** 5 years, 1 month ago

Read this:

To avoid API Management from making this HTTP request again, when the same user makes another request, you can specify to store the user profile in the cache.

XML

Copy

<cache-store-value

key="@("userprofile-" + context.Variables["enduserid"])"

value="@((string)context.Variables["userprofile"])" duration="100000" />

API Management stores the value in the cache using the exact same key that API Management originally attempted to retrieve it with. The duration that API Management chooses to store the value should be based on how often the information changes and how tolerant users are to out-of-date information.

It is important to realize that retrieving from the cache is still an out-of-process, network request and potentially can still add tens of milliseconds to the request. The benefits come when determining the user profile information takes longer than that due to needing to do database queries or aggregate information from multiple back-ends.

The final step in the process is to update the returned response with the user profile information.

upvoted 3 times

  **LEOPOLD** 5 years, 1 month ago



update the returned response = outbound

upvoted 4 times

  **oxaytol** 4 years, 11 months ago

it is outbound as balddie sent in the following url <https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-cache>

upvoted 2 times

  **balddie** 5 years, 4 months ago

I think its outbound - <https://docs.microsoft.com/en-us/azure/api-management/api-management-howto-cache>

upvoted 5 times

🗨️ 👤 **dev\_nathan\_sg** Most Recent 4 years, 6 months ago

The value is only available after processing, so the storing of the value has to be Outbound, therefore the answer of Inbound, Inbound, Outbound, Outbound is correct.

upvoted 5 times

🗨️ 👤 **Santosh564** 4 years, 7 months ago

<outbound>

<cache-store duration="seconds" />

<base />

</outbound>

upvoted 4 times

🗨️ 👤 **Satheesh88** 4 years, 5 months ago

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

upvoted 1 times

🗨️ 👤 **MrMohan** 4 years, 9 months ago

cache-store value should be outbound

upvoted 3 times

🗨️ 👤 **UomoLumaca** 4 years, 9 months ago

I think the wording is ambiguous in this one.

Do I need to cache the individual API's response or the user's profile? The question says both things, which can contradict each other.

If it is the first one, the cache store value must be outbound (I cannot know what to store if I don't have the API's response - yes, I could call the individual API in the inbound policy but then it would mean that I would never use the real underlying API call in any case, and I don't think that's intended).

Otherwise, the procedure in the link rrongcheng provided must be followed (which makes a call to the user profile data API in the inbound policy).

I would go with the second one (so, inbound) since it references an existing procedure in Microsoft's documentation.

upvoted 1 times

🗨️ 👤 **Steve\_az** 4 years, 9 months ago

Options - <https://www.examtopycs.com/assets/media/exam-media/03135/0006700001.jpg>

upvoted 11 times

🗨️ 👤 **hertino** 4 years, 9 months ago

Today I got 900, I had this question, my answer:

A set-variable Inbound.

cache-lookup-value Inbound

cache-store-value Inbound

Find-and-replace Outbound

upvoted 18 times

🗨️ 👤 **sinh** 4 years, 7 months ago

Even a score of 900 doesn't necessarily mean your answer is correct.

upvoted 9 times

🗨️ 👤 **bhushan\_786** 4 years, 5 months ago

He is just guiding us to the most possible and correct answer. We know it can be wrong also. If you are not agreeing with him, ignore it. No need of taking a dig at him. :)

upvoted 11 times

🗨️ 👤 **rrongcheng** 4 years, 10 months ago

The answer should be:

1. set-variable -> Inbound

2. cache-lookup-value -> Inbound

3. cache-store-value -> inbound

4. find-and-replace -> outbound

Example from doc:

"Once you combine all these steps together, the end result is a policy that looks like the following one."

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key#:~:text=To%20avoid%20API%20Management%20from,user%20profile%20in%20the%20cache.&text=API%20Management%20stores%20the%20value,atte>

upvoted 6 times

🗨️ 👤 **anishk** 5 years ago




Here the user ID of the client must be detected and then the response must be cached for the given user ID. vs General caching output for all users.

So to detect userid, its inbound . Check code on the following site (bottom of the page) for this scenario

<https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

upvoted 2 times

  **Dumindu** 5 years, 2 months ago

given ans is correct



upvoted 7 times

  **gssicftohbcrcsgtjg** 5 years, 3 months ago

options

<https://encrypted-tbn0.gstatic.com/images?q=tbn%3AAND9GcQa0WK1UePzrGPAJgWxnZJP5tVh4aRossPenlrzT4AXtipbsnT2>

upvoted 15 times



  **JohnyT** 5 years, 4 months ago

cache-store-value is inbound. It must have corresponding cache-lookup.

cache-store is outbound. It only sets the max-age header.

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

upvoted 7 times

  **Coder1** 4 years, 9 months ago

Indeed its Inbound, link : <https://docs.microsoft.com/en-us/azure/api-management/api-management-sample-cache-by-key>

upvoted 1 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:

**Suggested Answer:** *Explanation*

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 be an error in the question. Should have enable instead of create.

Box 5: all -

Encrypt both data and operating system.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/virtual-machines/linux/encrypt-disks>

  **031920** Highly Voted  4 years, 9 months ago

<https://vceguide.com/how-should-you-complete-the-azure-cli-commands-2/>

upvoted 8 times

  **Kratos** Most Recent  5 years ago

Reference: <https://docs.microsoft.com/en-us/azure/virtual-machines/linux/disk-encryption-cli-quickstart>

upvoted 2 times

  **VMCoder** 5 years, 2 months ago

Options plz

upvoted 3 times

## HOTSPOT -

You are developing an Azure App Service hosted ASP.NET Core web app to deliver video on-demand streaming media. You enable an Azure Content Delivery

Network (CDN) Standard for the web endpoint. Customer videos are downloaded from the web app by using the following example

URL: <http://www.contoso.com/content.p4?quality=1>

All media content must expire from the cache after one hour. Customer videos with varying quality must be delivered to the closest regional point of presence

(POP) node.

You need to configure Azure CDN caching rules.

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

NOTE: Each correct selection is worth one point.

Hot Area:

**Suggested Answer:** *Explanation*

Box 1: Override -

Override: Ignore origin-provided cache duration; use the provided cache duration instead. This will not override cache-control: no-cache.

Set if missing: Honor origin-provided cache-directive headers, if they exist; otherwise, use the provided cache duration.

Incorrect:

Bypass cache: Do not cache and ignore origin-provided cache-directive headers.

Box 2: 1 hour -

All media content must expire from the cache after one hour.

Box 3: Cache every unique URL -

Cache every unique URL: In this mode, each request with a unique URL, including the query string, is treated as a unique asset with its own cache. For example, the response from the origin server for a request for example.ashx?q=test1 is cached at the POP node and returned for subsequent caches with the same query string. A request for example.ashx?q=test2 is cached as a separate asset with its own time-to-live setting.




Incorrect Answers:

Bypass caching for query strings: In this mode, requests with query strings are not cached at the CDN POP node. The POP node retrieves the asset directly from the origin server and passes it to the requestor with each request.

Ignore query strings: Default mode. In this mode, the CDN point-of-presence (POP) node passes the query strings from the requestor to the origin server on the first request and caches the asset. All subsequent requests for the asset that are served from the POP ignore the query strings until the cached asset expires.

References:

<https://docs.microsoft.com/en-us/azure/cdn/cdn-query-string>

  **czmiel24**  5 years, 3 months ago


<https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&ved=2ahUKEwjBj97okOXnAhV8w8QBHY-OBT4QFjAAegQIBRAB&url=https%3A%2F%2Fvceplus.com%2Fexams%2F4477%2Faz-203%2F17990%2Fmicrosoft-actualtests-az-203-v2019-06-16-by-linda-46q.pdf&usg=AOvVaw1VZbMqtl6yVly5YQlrvsp>

upvoted 27 times

  **Balajicr7** 4 years, 9 months ago

Not all superhero wear capes some are named as czmiel24


upvoted 18 times

  **ozigee**  5 years ago

Options. Sry, with answers:

<https://html2-f.scribdassets.com/8n8ya0pvi87d2upg/images/10-33baf3d361.jpg>

upvoted 15 times

  **oxaytol** 4 years, 11 months ago

Thanks for the info.

upvoted 1 times

☐  **gunencali** Most Recent 4 years, 5 months ago


Answer is correct.

Override,

1 hour,

cache every unique URL

upvoted 3 times

☐  **27close** 4 years, 5 months ago

Override, 1 hour, cache every unique URL

upvoted 3 times

☐  **NPoul\_92** 4 years, 8 months ago

<https://vceguide.com/which-options-should-you-use-3/>

upvoted 4 times

You have an Azure App Services Web App Azure SQL Database instance Azure Storage Account, and an Azure Redis Cache instance in a resource group.

A developer must be able to publish code to the web app. You must grant the developer the Contributor role to the web app.

You need to grant the role.

Which two commands can you use? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. az role assignment create
- B. az role definition create
- C. New-AzureRmRoleAssignment
- D. New-AzureRmRoleDefinition

**Suggested Answer: AC**

A: The az role assignment create command creates a new role assignment for a user, group, or service principal.

Example: Create role assignment for an assignee.

az role assignment create --assignee sp\_name --role a\_role

C: The New-AzureRmRoleAssignment command assigns the specified RBAC role to the specified principal, at the specified scope.

Incorrect Answers:

B, D: Creates a custom role in Azure RBAC.

References:

<https://docs.microsoft.com/en-us/cli/azure/role/assignment?view=azure-cli-latest#az-role-assignment-create> <https://docs.microsoft.com/en-us/powershell/module/azurerm.resources/new-azureroleassignment?view=azurermps-6.13.0>



  **hertino** Highly Voted 4 years, 9 months ago

Today I got 900, I had this question, my answer:

az role assignment create

New-AzureRmRoleAssignment

upvoted 25 times

  **MyLife** 4 years, 6 months ago


Answer with justification is more appropriate than your marks here. No body cares.

upvoted 12 times

  **NajamKhan** 4 years, 6 months ago

Even getting 900 doesn't mean, your answer is correct.

upvoted 9 times

  **SidSadhu** Highly Voted 4 years, 9 months ago

Correct answer

upvoted 5 times

  **exam\_taker\_1** Most Recent 1 year, 4 months ago

Answer is correct. You want to assign existing role, not create a new custom role

upvoted 1 times

  **SidSadhu** 4 years, 9 months 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. OAuth Client Credential Grant
- B. Basic Authentication
- C. Certificate Authentication
- D. Digest Authentication

**Suggested Answer: AC**

  **stefanat**  5 years, 3 months ago

Should be BC

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

upvoted 47 times

  **lau13** 5 years ago

For those who are asking is it AC or BC

The requirement is asking for the backend service to include a valid HTTP authorization header, B is the minimum that we can do to achieve the goal; you can do A, but it is overkill and it will require more steps to achieve, not the best option for this requirement.

upvoted 2 times

  **Juanlu** 4 years, 5 months ago

Agree. As @lq said, option A is not a policy Microsoft offers

upvoted 1 times

  **shuakwe**  5 years ago

It should be AB.

Certification authentication does not use the Authorization header so it doesn't match the requirements.

Example of implementing OAuth Client Credential Grant in API Management: <https://platform.deloitte.com.au/articles/oauth2-client-credentials-flow-on-azure-api-management>

upvoted 13 times

  **dailak** 4 years, 8 months ago

I think AB is correct because Basic and OAuth are using Authorization Header

upvoted 2 times

  **altafpatel1984** 3 years, 6 months ago

I agree

upvoted 1 times

  **hems4all**  3 years, 11 months ago

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.

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

upvoted 1 times

🗨️ 👤 **SidSadhu** 4 years, 9 months ago

B, C

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

upvoted 2 times

🗨️ 👤 **bugimachi** 4 years, 4 months ago

This is true!

For all who make the same mistake as I did (because it did not read carefully enough): The question is about authentication between APIM and backend (and has nothing to do with the user authentication!)

upvoted 2 times

🗨️ 👤 **AK89** 5 years ago

Correct answer are B and C

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

upvoted 10 times

🗨️ 👤 **Iq** 5 years ago

Should be B and C

option A is not a policy Microsoft offers

upvoted 9 times

🗨️ 👤 **Iloves73g** 5 years ago

B,C correct

upvoted 11 times

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

AC is correct

upvoted 6 times

🗨️ 👤 **heero** 5 years, 3 months ago

So it should be AC or BC ?

upvoted 1 times



## DRAG DROP -

You maintain an existing Azure SQL Database instance. Management of the database is performed by an external party. All cryptographic keys are stored in an

Azure Key Vault.

You must ensure that the external party cannot access the data in the SSN column of the Person table.

Will each protection method meet the requirement? To answer, drag the appropriate responses to the correct protection methods. Each response 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:

**Suggested Answer:** *Explanation*

Box 1: Yes -

You can configure Always Encrypted for individual database columns containing your sensitive data. When setting up encryption for a column, you specify the information about the encryption algorithm and cryptographic keys used to protect the data in the column.

Box 2: No -

Box 3: Yes -



In SQL Database, the VIEW permissions are not granted by default to the public fixed database role. This enables certain existing, legacy tools (using older versions of DacFx) to work properly. Consequently, to work with encrypted columns (even if not decrypting them) a database administrator must explicitly grant the two VIEW permissions.

Box 4: No -

All cryptographic keys are stored in an Azure Key Vault.

References:



<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine>

  **Daltonic75** Highly Voted 5 years, 2 months ago

Same question in <https://www.examttopics.com/exams/microsoft/az-300/view/31/>

But different answers. Check discussion in that URL.

upvoted 14 times

  **babu789** 4 years, 10 months ago

@daltonic75 you are awesome.

upvoted 3 times

  **saeza100** Highly Voted 4 years, 11 months ago



All of them "No", we have "always Encrypt" but haven't "AlwaysOn Encrypt"

upvoted 8 times

  **Juanlu** 4 years, 5 months ago



Agree. NO | NO | NO | NO

upvoted 2 times

  **Ananas** 4 years, 9 months ago

Actually the solution talks about "always encrypt", that is why it is marked as "yes"

upvoted 3 times

  **shrusudev** Most Recent 4 years, 1 month ago

<https://vceguide.com/drag-drop-666/>

upvoted 1 times

  **sumaiyap86** 4 years, 2 months ago

Box 4 is NO

The Database Engine stores encryption configuration for each column in database metadata. Note, however, the Database Engine never stores or uses the keys of either type in plaintext. It only stores encrypted values of column encryption keys and the information about the location of column master keys, which are stored in external trusted key stores, such as Azure Key Vault, Windows Certificate Store on a client machine, or a hardware security module.

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine?view=sql-server-ver15>  
upvoted 1 times

🗨️ 👤 **cbn** 4 years, 3 months ago

I think the last option should be YES.

"Note, however, the Database Engine never stores or uses the keys of either type in plaintext. It only stores encrypted values of column encryption keys and the information about the location of column master keys, which are stored in external trusted key stores, such as Azure Key Vault"

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine?view=sql-server-ver15>  
upvoted 1 times

🗨️ 👤 **cbn** 4 years, 3 months ago

Therefore, YES | NO | YES | YES

upvoted 1 times

🗨️ 👤 **funfun** 4 years, 4 months ago

box 3 should be yes, according to MS, public role do not have VIEW permission by default

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/always-encrypted-database-engine?view=sql-server-ver15>

upvoted 2 times

🗨️ 👤 **Meenusiju** 4 years, 7 months ago

Box 3 is "NO" as per <https://www.examtopycs.com/exams/microsoft/az-300/view/31/> discussions.

upvoted 2 times

🗨️ 👤 **Daltonic75** 5 years, 2 months ago

Options (same question): <https://www.examtopycs.com/exams/microsoft/az-300/view/31/>

Or: <https://www.examtopycs.com/assets/media/exam-media/02758/0043800001.jpg>

upvoted 5 times

You develop an Azure web app. You monitor performance of the web app by using Application Insights. You need to ensure the cost for Application Insights does not exceed a preset budget. What should you do?

- A. Implement ingestions sampling using the Application Insights SDK.
- B. Set a daily cap for the Application Insights instance.
- C. Implement ingestion sampling using the Azure portal.
- D. Implement adaptive sampling using the Azure portal.
- E. Implement adaptive sampling using the Application Insights SDK.

**Suggested Answer: E**

Sampling is an effective way to reduce charges and stay within your monthly quota.

You can set sampling manually, either in the portal on the Usage and estimated costs page; or in the ASP.NET SDK in the .config file; or in the Java SDK in the

ApplicationInsights.xml file, to also reduce the network traffic.

Adaptive sampling is the default for the ASP.NET SDK. Adaptive sampling automatically adjusts to the volume of telemetry that your app sends. It operates automatically in the SDK in your web app so that telemetry traffic on the network is reduced.




Incorrect Answers:

B: You can use the daily volume cap to limit the data collected.

To change the daily cap, in the Configure section of your Application Insights resource, in the Usage and estimated costs pane, select Daily Cap.

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling>

  **Regimiento**  5 years, 5 months ago

The answer should be B:

setting daily cap will prevent exceeding predefined limit.

E is incorrect, as it only reduces the amount of data, but not limits it to specific max amount.

upvoted 35 times

  **cbn** 4 years, 3 months ago

Agree. Website is 'being developed' and need to monitor performance while not exceeding the preset budget.

upvoted 1 times

  **KK\_uniq**  5 years, 3 months ago

Sorry it will be E only.

You can use the daily volume cap to limit the data collected. However, if the cap is met, a loss of all telemetry sent from your application for the remainder of the day occurs. It is not advisable to have your application hit the daily cap. You can't track the health and performance of your application after it reaches the daily cap.

Instead of using the daily volume cap, use sampling to tune the data volume to the level you want. Then, use the daily cap only as a "last resort" in case your application unexpectedly begins to send much higher volumes of telemetry.

upvoted 16 times

  **pieixoto** 5 years ago

While your answer is best practice it does not meet the requirement. The requirement is to "ensure the cost for Application Insights does not exceed a preset budget." You cannot do this with adaptive sampling alone.

upvoted 5 times

  **[Removed]** 5 years ago

B & E seem right to me, together. Otherwise only B. E is good practice but doesn't keep the cost down which is what the question is asking.

upvoted 2 times

  **aftab7500** 4 years, 7 months ago

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/pricing>

upvoted 2 times

🗨️ **cbn** 4 years, 3 months ago

Website is 'being developed' and need to monitor performance while not exceeding the preset budget.

I will go with "B"  
upvoted 1 times

🗨️ **AshwaniT** Most Recent 4 years, 2 months ago

'B' is the correct answer here.  
upvoted 1 times

🗨️ **ss22** 4 years, 3 months ago

Answer should be B  
upvoted 1 times

🗨️ **priyankareddy** 4 years, 5 months ago

Reason for choosing B.

<https://docs.microsoft.com/en-us/azure/azure-monitor/platform/manage-cost-storage#understand-your-usage-and-estimate-costs>.

The daily cap is intended to be used as a way to manage an unexpected increase in data volume from your managed resources and stay within your limit, or when you want to limit unplanned charges for your workspace.

upvoted 2 times

🗨️ **Cornholioz** 4 years, 5 months ago

Once Daily cap is reached, it will stop pumping info. Loss.

Ingestion sampling is for server side.

Adaptive Sampling seems right. To be done using SDK.

I'll go with E.

upvoted 1 times

🗨️ **novice25** 4 years, 7 months ago

Why not D an option here, pls suggest?  
upvoted 1 times

🗨️ **rcisz** 4 years, 6 months ago

I think D. You can implement after observation of volume of metrics, how to set sampling .. as stated in documentation  
upvoted 1 times

🗨️ **MAzureLearner2020** 4 years, 7 months ago

There is no information about the web app.

Adaptive sampling is only available for ASP.NET applications.

Ingestion sampling available if no other sampling is in effect.

However I guess it should be B, since it is the only way to ensure it.

upvoted 3 times

🗨️ **Larry616** 4 years, 9 months ago

Why not A? A can control the quota

Ingestion sampling happens at the Application Insights service endpoint. It discards some of the telemetry that arrives from your app, at a sampling rate that you set. It doesn't reduce telemetry traffic sent from your app, but helps you keep within your monthly quota. The main advantage of ingestion sampling is that you can set the sampling rate without redeploying your app. Ingestion sampling works uniformly for all servers and clients, but it does not apply when any other types of sampling are in operation.

upvoted 1 times

🗨️ **Larry616** 4 years, 9 months ago

Sorry, A is wrong coz it happens on server side, only can be configured in portal rather than SDK on client side.

upvoted 2 times

🗨️ **harry973** 4 years, 9 months ago

With the ASP.NET SDK's adaptive sampling, the data volume is adjusted automatically to keep within a specified maximum rate of traffic for default Application Insights monitoring. If the application produces a low amount of telemetry, such as when debugging or due to low usage, items won't be dropped by the sampling processor as long as volume is below the configured events per second level. For a high volume application, with the default threshold of five events per second, adaptive sampling will limit the number of daily events to 432,000. Using a typical average event size of 1 KB, this corresponds to 13.4 GB of telemetry per 31-day month per node hosting your application (since the sampling is done local to each node.)

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/pricing>

so, cost limit can be decided using adaptive sampling also. Also, the questions says we need to monitor the web app which wont be possible if daily cap is met.

upvoted 1 times

🗳️ 👤 **jw\_duke** 4 years, 10 months ago

I would go with E due to the operational impact when setting a cap on the Application Insight.

During the high traffic volume, the cap limit would be easily reached and application insight would stop collecting data.

upvoted 1 times

🗳️ 👤 **brajen** 4 years, 10 months ago

Answer is B:

You can set a daily cap to ensure that the cost does not exceed a budget. You can go to "Usage and estimated costs" for your Application Insights resource and then hit on Daily cap.

upvoted 1 times

🗳️ 👤 **pranay** 4 years, 10 months ago

The main advantages of sampling are:

1.Application Insights service drops ("throttles") data points when your app sends a very high rate of telemetry in a short time interval. Sampling reduces the likelihood that your 2.application will see throttling occur.

3.To keep within the quota of data points for your pricing tier.

4.To reduce network traffic from the collection of telemetry.

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/sampling>

As per 3rd point sampling has to be used to keep data withing limit and preferred sampling is always adaptive. So answer should be E

upvoted 1 times

🗳️ 👤 **oxaytol** 4 years, 11 months ago

This one is tricky, i'm pretty sure that given answer is correct, since it's described in best practices as follows:

"Sampling is a feature in Azure Application Insights. It is the recommended way to reduce telemetry traffic, data costs, and storage costs, while preserving a statistically correct analysis of application data"

upvoted 2 times

🗳️ 👤 **AK89** 5 years ago

The answer should be B

upvoted 2 times

🗳️ 👤 **riley5** 5 years, 1 month ago

"Adaptive sampling is enabled by default in all the latest versions of the Application Insights ASP.NET and ASP.NET Core Software Development Kits (SDKs). It is also used by Azure Functions."

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/pricing>

If its enabled by default, then it's already implemented, which begs the question if the default answer is correct for this question.

upvoted 1 times

🗳️ 👤 **riley5** 5 years, 1 month ago

here are the steps given in the documentation for dealing with the issue:

A. Review your data volume for the month. This includes all the data that's received and retained (after any sampling) from your server and client apps, and from availability tests.

B. A separate charge is made for multi-step web tests. (This doesn't include simple availability tests, which are included in the data volume charge.)

C. View data volume trends for the past month.

D. Enable data ingestion sampling.

E. Set the daily data volume cap.

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/pricing>

From this, both ingestion sampling and setting a daily cap volume seem to be correct.

upvoted 1 times

🗳️ 👤 **Ananas** 4 years, 9 months ago

From this, they seem to be correct when used together

upvoted 1 times

🗳️ 👤 **realbart** 5 years, 1 month ago

## Types of sampling



There are three different sampling methods:

Adaptive sampling automatically adjusts the volume of telemetry sent from the SDK in your ASP.NET/ASP.NET Core app, and from Azure Functions. This is the default sampling when you use the ASP.NET or ASP.NET Core SDK. Adaptive sampling is currently only available for ASP.NET server-side telemetry, and for Azure Functions.

Fixed-rate sampling reduces the volume of telemetry sent from both your ASP.NET or ASP.NET Core or Java server and from your users' browsers. You set the rate. The client and server will synchronize their sampling so that, in Search, you can navigate between related page views and requests.

Ingestion sampling happens at the Application Insights service endpoint. It discards some of the telemetry that arrives from your app, at a sampling rate that you set. It doesn't reduce telemetry traffic sent from your app, but helps you keep within your monthly quota. The main advantage of ingestion sampling is that you can set the sampling rate without redeploying your app. Ingestion sampling works uniformly for all servers and clients, but it does not apply when any other types of sampling are in operation.

upvoted 5 times

  **riley5** 5 years, 1 month ago

From this, are you saying ingestion sampling is correct?

upvoted 2 times

You are developing an ASP.NET Core Web API web service. The web service uses Azure Application Insights for all telemetry and dependency tracking. The web service reads and writes data to a database other than Microsoft SQL Server.

You need to ensure that dependency tracking works for calls to the third-party database.

Which two Dependency Telemetry properties should you store in the database? Each correct answer presents part of the solution.



NOTE: Each correct selection is worth one point.

- A. Telemetry.Context.Operation.Id
- B. Telemetry.Name
- C. Telemetry.Context.Cloud.RoleInstance
- D. Telemetry.Context.Session.Id
- E. Telemetry.Id

**Suggested Answer:** AE

References:

<https://docs.microsoft.com/en-us/azure/azure-monitor/app/custom-operations-tracking>

 **AK89**  5 years ago

Correct : Operations.ID and Request.ID/Telemetry.ID

upvoted 6 times

You are creating a hazard notification system that has a single signaling server which triggers audio and visual alarms to start and stop. You implement Azure Service Bus to publish alarms. Each alarm controller uses Azure Service Bus to receive alarm signals as part of a transaction. Alarm events must be recorded for audit purposes. Each transaction record must include information about the alarm type that was activated.

You need to implement a reply trail auditing solution. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Assign the value of the hazard message MessageId property to the DeliveryCount property.
- B. Assign the value of the hazard message SequenceNumber property to the DeliveryCount property.
- C. Assign the value of the hazard message MessageId property to the SequenceNumber property.
- D. Assign the value of the hazard message MessageId property to the CorrelationId property.
- E. Assign the value of the hazard message SessionId property to the SequenceNumber property..
- F. Assign the value of the hazard message SessionId property to the ReplyToSessionId property.

**Suggested Answer:** DF

D: CorrelationId: Enables an application to specify a context for the message for the purposes of correlation; for example, reflecting the MessageId of a message that is being replied to.

F: ReplyToSessionId: This value augments the ReplyTo information and specifies which SessionId should be set for the reply when sent to the reply entity.

Incorrect Answers:

A, B: DeliveryCount -


Number of deliveries that have been attempted for this message. The count is incremented when a message lock expires, or the message is explicitly abandoned by the receiver. This property is read-only.

C, E: SequenceNumber -

The sequence number is a unique 64-bit integer assigned to a message as it is accepted and stored by the broker and functions as its true identifier. For partitioned entities, the topmost 16 bits reflect the partition identifier. Sequence numbers monotonically increase and are gapless. They roll over to 0 when the 48-64 bit range is exhausted. This property is read-only.

References:

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messages-payloads>

 **Cornholioz** Highly Voted 4 years, 5 months ago

Given answer correct.

CorrelationId and ReplyToSessionId are meant for this purpose. To correlate and to record SessionId value in the reply. The rest them (DeliveryCount and SequenceNumber) are reserved for their own purposes like the names suggest and are hence made read-only too.

upvoted 6 times

 **HexWr** Highly Voted 4 years, 5 months ago

As per the docs, DeliveryCount and SequencyNumber are read-only, so the given answer must be correct

upvoted 5 times

 **cbn** Most Recent 4 years, 3 months ago

Given answer is correct.

However, option F should have been the other way round like below:

Assign the value of the hazard message "ReplyToSessionId" property to the "SessionID" property.

upvoted 1 times

 **santosh\_rajp** 4 years, 4 months ago

Answers are correct.

upvoted 2 times

 **JPsukum** 4 years, 7 months ago



why not B & E is correct answer ? Could see B & E in few other question sets..

upvoted 1 times

You develop a gateway solution for a public facing news API. The news API back end is implemented as a RESTful service and uses an OpenAPI specification.

You need to ensure that you can access the news API by using an Azure API Management service instance.

Which Azure PowerShell command should you run?

- A. `Import-AzureRmApiManagementApi ""Context $ApiMgmtContext ""SpecificationFormat "Swagger" -SpecificationPath $SwaggerPath ""Path $Path`
- B. `New-AzureRmApiManagementBackend -Context $ApiMgmtContext -Url $Url -Protocol http`
- C. `New-AzureRmApiManagement ""ResourceGroupName $ResourceGroup ""Name $Name ""Location $Location ""Organization $Org ""AdminEmail $AdminEmail`
- D. `New-AzureRmApiManagementBackendProxy ""Url $ApiUrl`

**Suggested Answer: D**

`New-AzureRmApiManagementBackendProxy` creates a new Backend Proxy Object which can be piped when creating a new Backend entity.

Example: Create a Backend Proxy In-Memory Object

```
PS C:\>$secpassword = ConvertTo-SecureString "PlainTextPassword" -AsPlainText -Force
```

```
PS C:\>$proxyCreds = New-Object System.Management.Automation.PSCredential ("foo", $secpassword)
```

```
PS C:\>$credential = New-AzureRmApiManagementBackendProxy -Url "http://12.168.1.1:8080" -ProxyCredential $proxyCreds
```

```
PS C:\>$apimContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
```

```
PS C:\>$backend = New-AzureRmApiManagementBackend -Context $apimContext -BackendId 123 -Url 'https://contoso.com/awesomeapi' -
Protocol http -Title
```

```
"first backend" -SkipCertificateChainValidation $true -Proxy $credential -Description "backend with proxy server"
```

Creates a Backend Proxy Object and sets up Backend

Incorrect Answers:

A: The `Import-AzureRmApiManagementApi` cmdlet imports an Azure API Management API from a file or a URL in Web Application Description Language (WADL),

Web Services Description Language (WSDL), or Swagger format.

B: `New-AzureRmApiManagementBackend` creates a new backend entity in Api Management.

C: The `New-AzureRmApiManagement` cmdlet creates an API Management deployment in Azure API Management.

References:

<https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/new-azurermapi-management-backend-proxy?view=azurerm-ps-6.13.0>

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

Shouldn't the correct answer be A? I don't see why it should be D like the comment says.

upvoted 10 times

  **boldarev** 5 years, 1 month ago

I think also, Correct answer is A. An OpenAPI Specification, formerly known as the Swagger Specification.

upvoted 5 times

  **MAzureLearner2020** 4 years, 7 months ago

yes should be a

```
PS C:\>$ApiMgmtContext = New-AzureRmApiManagementContext -ResourceGroupName "Api-Default-WestUS" -ServiceName "contoso"
```

```
PS C:\>Import-AzureRmApiManagementApi -Context $ApiMgmtContext -SpecificationFormat "Swagger" -SpecificationPath
```



```
"C:\contoso\specifications\echoapi.swagger" -Path "apis"
```

upvoted 1 times

  **Juanlu** 4 years, 5 months ago

Yes, in my opinion should be A, but it look like is a little confuse.



upvoted 1 times

  **AK89** Highly Voted 5 years ago

Answer is C


Link : <https://docs.microsoft.com/en-us/azure/api-management/powershell-create-service-instance#create-an-api-management-service>

upvoted 10 times

  **Kumar1988** Most Recent 4 years, 2 months ago



should be D correct ans

upvoted 1 times

  **dancsita** 4 years, 3 months ago

According to Whizlabs the Answer is C

upvoted 2 times

  **dt dame** 4 years, 3 months ago

New-AzureRmApiManagement is used to create a new instance of the service if I am not wrong about this, then B C D are not relevant to me, A looks like the best option !

upvoted 1 times

  **ahadjithoma** 4 years, 5 months ago

You need another Api management service instance for testing. This is done with New-AzApiManagement. So it should be C



<https://docs.microsoft.com/en-us/azure/api-management/powershell-create-service-instance>

upvoted 6 times

  **bhushan\_786** 4 years, 5 months ago

Yes the answer is C.

upvoted 3 times

  **27close** 4 years, 5 months ago

answer c New-AzApiManagement -Name "myapim" -ResourceGroupName "myResourceGroup" -Location "West US" -Organization "Contoso" -AdminEmail "admin@contoso.com"

<https://docs.microsoft.com/en-us/azure/api-management/powershell-create-service-instance#create-an-api-management-service>



upvoted 2 times

  **NajamKhan** 4 years, 6 months ago

The C is the the correct answer.

Here is the answer: <https://docs.microsoft.com/en-us/azure/api-management/powershell-create-service-instance#create-an-api-management-service>

upvoted 2 times

  **AnkitN** 4 years, 9 months ago

Option A is incorrect since this is used to import an Azure API Management API from a file or a URL in Web Application Description Language (WADL), Web Services Description Language (WSDL), or Swagger format.

Option B is incorrect since this is used to create a new backend for the API.

Option D is incorrect since this is used is just used to create a new backend proxy.

upvoted 2 times

  **cbn** 4 years, 3 months ago

Agree. We need a new API management instance, so answer is C

upvoted 1 times

  **TangAnna** 3 years ago

No, we need to import API as A did

upvoted 1 times

  **ens1z** 4 years, 12 months ago

I think A is the best.

<https://docs.microsoft.com/en-us/powershell/module/azurerm.apimanagement/import-azurermapi?view=azurerm-6.13.0>

upvoted 2 times

  **lau13** 5 years ago

Import-AzApiManagementApi will import the WSDL/swagger of the API to the application, it cannot help you to ensure you have access to that API.

The question is saying - you are developing a gateway solution that will consume a public news API, i.e. that news API backend already exist, you do not need to run New-AzureRmApiManagementBackend to create it again.



That's why A & C are both incorrect.

upvoted 3 times

  **lau13** 5 years ago

typo I mean what's why A & B are both incorrect

upvoted 2 times

  **barubs** 5 years ago

It might be A because:

<https://docs.microsoft.com/en-us/azure/api-management/scripts/powershell-import-api-and-add-to-product?toc=/powershell/module/toc.json>

# import api from Url

But they should clarify what they are asking for...

upvoted 2 times

  **rrongcheng** 4 years, 10 months ago

Yes, I think it's A.

The question is to import the existing API to an API Management product. So the cmd is:

# import api from Url

\$api = Import-AzApiManagementApi -Context \$context -SpecificationUrl \$swaggerUrl -SpecificationFormat Swagger -Path \$apiPath

upvoted 1 times

  **Jay\_2pt0** 5 years, 1 month ago

Wouldn't the correct answer be C?

<https://docs.microsoft.com/en-us/azure/api-management/powershell-create-service-instance#create-an-api-management-service>

upvoted 6 times

  **hmirko** 5 years, 1 month ago

Can someone provide link with explanation why it should be Backend Proxy and not Backend? To me, it looks like answer is Backend Proxy, since PowerShell for Backend has parameter -Protocol http, which is not well for public facing news API

upvoted 2 times

You are developing a project management service by using ASP.NET. The service hosts conversations, files, to-do lists, and a calendar that users can interact with at any time.

The application uses Azure Search for allowing users to search for keywords in the project data.

You need to implement code that creates the object which is used to create indexes in the Azure Search service.

Which two objects should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. SearchCredentials
- B. SearchIndexClient
- C. SearchServiceClient
- D. SearchService

**Suggested Answer: BC**

The various client libraries define classes like Index, Field, and Document, as well as operations like Indexes.Create and Documents.Search on the

SearchServiceClient and SearchIndexClient classes.

Example:

The sample application we'll be exploring creates a new index named "hotels", populates it with a few documents, then executes some search queries. Here is the main program, showing the overall flow:

/ This sample shows how to delete, create, upload documents and query an index static void Main(string[] args)

```
{
IConfigurationBuilder builder = new ConfigurationBuilder().AddJsonFile("appsettings.json");
IConfigurationRoot configuration = builder.Build();
SearchServiceClient serviceClient = CreateSearchServiceClient(configuration);
Console.WriteLine("{0}", "Deleting index...\n");
DeleteHotelsIndexIfExists(serviceClient);
Console.WriteLine("{0}", "Creating index...\n");
CreateHotelsIndex(serviceClient);
ISearchIndexClient indexClient = serviceClient.Indexes.GetClient("hotels");
```

References:

<https://docs.microsoft.com/en-us/azure/search/search-howto-dotnet-sdk>

  **lywwind**  5 years, 4 months ago

Correct Answer: A C ?

SearchServiceClient serviceClient = new SearchServiceClient(searchServiceName, new SearchCredentials(adminApiKey));

upvoted 36 times

  **Dumindu** 5 years, 2 months ago

this is the correct answer

upvoted 4 times

  **Dumindu** 5 years, 2 months ago

A and C

upvoted 4 times

  **courtguard** 2 years, 2 months ago

SearchServiceClient(String, SearchCredentials) - Initializes a new instance of the SearchServiceClient class.

upvoted 1 times

  **[Removed]** 5 years ago

I think the answer BC is correct. The SearchServiceClient has many other constructors apart from the one using SearchCredentials object.

<https://docs.microsoft.com/en-us/dotnet/api/microsoft.azure.search.searchserviceclient?view=azure-dotnet#constructors>

But we will always need a combination of SearchServiceClient and SearchIndexClient to index.

upvoted 3 times